Rule Rationality

Yuval Heller[*]   Eyal Winter[†]

April 13, 2015

Abstract

We study the strategic advantages of following rules of thumb that bundle different games together (called *rule rationality*) when this may be observed by one's opponent. We present a model in which the strategic environment determines which kind of rule rationality is adopted by the players. We apply the model to characterize the induced rules and outcomes in various interesting environments. Finally, we show the close relations between act rationality and "Stackelberg stability" (no player can earn from playing first). JEL classification: C72, D82.

1 Introduction

Act rationality is the notion that perfectly rational agents choose actions that maximize their utility in every specific situation. However, both introspection and experimental evidence suggest that people behave differently and follow "rule rationality": they adopt rules of thumb, or modes of behavior, that maximize some measure of average utility taken over the set of decision situations to which that rule applies; then, when making a decision, they choose an action that accords with the rule they have adopted. As a result of bundling together many decision situations, this rule of thumb induces smaller cognitive costs and fewer informational requirements (see Baumol & Quandt, 1964; Harsanyi, 1977; Ellison & Fudenberg, 1993; Aumann, 2008, and the references there).

Reducing the cognitive costs is arguably the main advantage of such rules of thumb, but it is not the only advantage. Rule-rational behavior can also arise from religious, moral, or ideological rules, which are unrelated to cognitive limitations. Such rules allow players to commit to a certain behavior and by that affecting the behavior of others to their advantage. For example, the Jewish rules of a kosher diet serve as an important commitment device. As eating together is an important social event, such dietary restrictions serve as a credible commitment to refrain from major social interaction with non-Jews and hence a commitment to engage socially within the community. A rule that bans eating with non-Jews would have been too stringent and too costly and would not have been able to survive. Likewise, when a state or a military organization chooses to follow the Geneva conventions regarding prisoners of war, it establishes a commitment that binds the incentives of the enemy to fight to the "last drop of blood." Politeness rules such as "ladies first" facilitate coordination.

[*]Department of Economics and the Queen's College, University of Oxford. yuval.heller@economics.ox.ac.uk. URL: https://sites.google.com/site/yuval26/.

[†]Center for the Study of Rationality and Department of Economics, Hebrew University of Jerusalem. mseyal@mscc.huji.ac.il . URL: http://www.ma.huji.ac.il/~mseyal/. The author is grateful to the German Israeli Foundation for Scientific Research and to Google for their financial support. Both authors are grateful to the very useful comments of the editor and the referees, and to helpful discussion with the audiences at Queen Mary spring theory workshop and in seminars in Oxford and HUJI.

In all these as well as similar examples, bounded cognition is of little relevance, and emotions, moral standards, and social norms dictate the rule. The role of commitment in strategic situations has been extensively investigated since the seminal work of Schelling, 1960 (see, e.g., Bade et al., 2009; Renou, 2009 and the references within for recent papers in this vast literature). While commitment has always been studied in the context of a single game its relationship to rule behavior in environments of multiple games/situations has not been studied. We believe that this relationship can shed light on a variety of social rules with which complexity and cognitive limitation has little to do.

In this paper we focus on rule rational behavior in which players bundle games together while interacting with others who gradually learn the structure of bundling (which allows the bundling to serve as a commitment device.)[1] *The main contribution of this paper is threefold: (1) We present a tractable model in which the properties of the strategic environment determine which kind of rule rationality is adopted by the players; (2) we characterize the induced rules and outcomes in various interesting environments; and (3) we show the close relations between act rationality and "Stackelberg stability" (no player can earn from playing before his opponent), and relate it to the value of information.*

The extent to which the commitment to act according to a moral or ideological principle is credible depends on comparing the cost of adhering to the principle relative to the cost of losing the commitment benefits in the future. These are the kind of considerations we employ in our model to define rule rationality. The declaration of such principles can be viewed as a signal that facilitates the gradual learning about one's bundling strategy by others. To be sure, these signals are not fully revealing but the fact that these declarations are very often made and often taken seriously suggests that stressing ideological or moral principles is not merely cheap talk either. Roth et al. (1991) demonstrate that a commitment to a principle or a moral standard becomes a social norm that people act upon.[2] In the presence of some pre-play communication people who contemplate deviating from the social norm (i.e., choosing a different partition) will use deception, which might be detected by the opponent.[3]

Another form of rule rationality relevant to our context includes firms or organizations that adopt rigid policies applying to a broad class of circumstances without elaborating the rationale of each policy. If a craftsman rejects an offer to accept 50% of his proposed price for doing 50% of the job on the basis that it is "unprofessional" or when a chain store rejects a returned item after the refund deadline even under exceptional circumstances by simply arguing that "it is against our policy," they act within this form of rule rationality.

We now briefly present our model. An *environment* is a finite set of two-player normal-form games that share the same set of feasible actions.[4] The environment is endowed with a function that determines the probability of each game being played. Each environment induces the following two-stage *meta-game:* at stage 1, each player chooses a partition over the set of games; at stage 2, each player chooses an action after observing the element of his partition that includes the realized game and his opponent's partition. A *rule-rational equilibrium* of the environment is defined as a subgame-perfect equilibrium of the meta-game. We interpret such an equilibrium as a stable outcome of a dynamic process of social learning, in which the choices of partitions in the first stage (interpreted as rules of thumb that bundle games

---

[1] This "commitment" advantage of rule rationality might be of importance also in single-player games as a "self commitment" device with a multiple-selves interpretation.

[2] Roth et al. (1991) study the ultimatum game across different cultures. They find that responders always reject offers below a certain threshold, and that these thresholds are culture-dependent (for example, they are substantially lower in Israel and Japan than in Slovenia and the U.S.). Proposers seem to be familiar with the distribution of the threshold within their own culture.

[3] See, e.g., the findings of Israel et al. (2014) which suggest that watching 30 seconds of pre-play communication is enough to predict players' choices better than a random guess.

[4] Assuming a fixed set of action is without loss of generality; see Footnotes 7 and 10.

together) and the behaviors in the second stage (given the partitions) evolve.

A rule-rational equilibrium is *act-rational* if players always use the finest partition and play a Nash equilibrium in all games, and it is *non-act-rational* if at least one of the players does not best reply in at least one game. The observability of the opponent's partition is a key assumption in our model. If the opponent's partition is completely unobservable, then one can show that all environments only admit act-rational equilibria. Many of our results can be extended to a setup of partial observability, and this is discussed in Section 8.2.

The following motivating example demonstrates that a non-act-rational equilibrium can strictly Pareto-dominate the set of all act-rational equilibrium payoffs.

Example 1. Consider an environment in which nature chooses with equal probability one of the two games shown in Table 1: The Prisoner's Dilemma (G_1) and the Stag-Hunt Game (G_2).

Table 1: Prisoner's Dilemma and Stag Hunt

G_1 (Prisoner's Dilemma)				G_2 (Stag Hunt)		
	C	D			C	D
C	9,9	0,9.5		C	9,9	0,8.5
D	9.5,0	1,1		D	8.5,0	1,1

These two games are relatively "similar" in the sense that the maximal difference between the two payoff matrices is 1. Note that in the Prisoner's Dilemma, (D, D) is the unique Nash equilibrium, while the Stag-Hunt game admits two pure Nash equilibria: (D, D) and the Pareto-dominant (C, C). In any act-rational equilibrium of this environment, players play the unique Nash equilibrium (D, D) in the Prisoner's Dilemma game. As the maximal payoff in the Stag Hunt game is 9, each player can achieve an expected (ex-ante) payoff of at most 5 ($\frac{1+9}{2}$) in any act-rational equilibrium. Consider the following non-act-rational equilibrium. Both players choose the coarse partition (i.e., bundle together both games), play C on the equilibrium path, and play D (in both games) if any player deviated while choosing his partition. This equilibrium yields both players a payoff of 9. Note that the coarse partition in this example can be interpreted as if each player did not pay attention to small differences in the payoffs (say, up to one utility point).

Section 4 presents additional motivating examples. Example 2 shows a non-act-rational equilibrium in a trade interaction that implements an outcome that strictly Pareto-dominates the set of act-rational payoffs; the equilibrium relies on a partially coarse partition in which the players pay attention only to the total gain from trade. Example 3 presents an environment that admits only non-act-rational equilibria. Finally, Examples 5 and 6 demonstrate that non-act-rational equilibria can fit experimentally observed behavior in the "ultimatum game" and the "chain-store game".

In Section 5 we characterize rule-rational equilibria in various families of environments. An environment has a uniform best-response correspondence if all of its games have the same best-response correspondence (Morris & Ui, 2004). Note that such games share the same sets of Nash equilibria and correlated equilibria (Aumann, 1974). Our first results (Propositions 1-3) show that rule rationality in environments with a uniform best-response correspondence induces behavior that is consistent with the correlated equilibria of the underlying games; moreover, if such an environment is large enough, and each game has a small probability, then the correlated equilibrium of each underlying game can be approximated by a rule-rational equilibrium. Thus, our model presents a novel explanation of how correlated equilibria can arise from rule rationality (without mediators; see the discussion at the end of Section 5.1).

Next, we study environments in which in each game the sum of the payoffs is constant. We show that such constant-sum environments always admit an act-rational equilibrium. They may admit also non-act-rational equilibria, but these equilibria are limited in two ways: they must yield the same ex-ante expected payoff as the act-rational equilibrium, and they are not *robust,* i.e., they are not stable to a perturbation that yields an arbitrarily small probability that the choice of partition is not observed by the opponent. Example 9 shows that these properties need not hold if the games are VNM-equivalent to zero-sum games (Morris & Ui, 2004; Moulin & Vial, 1978).

Proposition 4 deals with environments in which each game admits a Pareto-dominant action profile, and we show that such environments admit an act-rational equilibrium that strictly Pareto-dominates all non-act-rational equilibria.[5] Finally, we generalize Example 1 in Section 5.4 and show that if an environment assigns high enough probability to game g^* that admits strict equilibrium a^* that dominates another strict equilibrium, then there is a "global" rule-rational equilibrium in which both players use the coarsest partition and play a^* in all the games.

In Section 6 we show the close relationship between the existence of act-rational equilibria and of Stackelberg stability. A Nash equilibrium in a game is *Stackelberg-stable* if it yields both players at least their *mixed-action Stackelberg payoff* (Mailath & Samuelson (2006)); that is, no player can guarantee a higher payoff by being a *Stackelberg leader*, i.e., playing first, and having his chosen strategy observed by his opponent before playing. A game is Stackelberg-stable if it admits a Stackelberg-stable equilibrium. Examples of Stackelberg-stable games are constant-sum games, games with a Pareto-dominant action profile, and games in which both players have dominant actions. We first show that if all games in the environment are Stackelberg-stable, then the environment admits an act-rational equilibrium. Second, we show that for any Stackelberg-unstable game G^*, there is an environment in which only G^* is Stackelberg-unstable, and the environment does not admit an act-rational equilibrium. Finally, we show that any act-rational equilibrium of a generic environment is preserved when a Stackelberg-stable game is added to the environment with small enough probability.

In Section 7 we extend our model in order to deal with a restricted set of feasible partitions. These restrictions may represent (1) exogenous incomplete information (a player is not able to observe some information about the opponent or nature), (2) salient information (a player can bundle together only sufficiently similar games), and (3) cultural factors (a player can categorize games in a way that is consistent with given social norms or moral concerns). We carry out this extension in order to analyze two applications: lemon markets (à la Akerlof (1970)) and partnership games. In both applications, we restrict the set of feasible partitions to those that bundle together situations only if the difference between them is sufficiently small (say, up to d). In the lemon market, we show that for any $d > 0$ the environment only admits non-act-rational equilibria with a positive level of trade (which increases continuously in d). In the partnership games, there is always a non-efficient act-rational equilibrium, while with a modest level of d the environment also admits an efficient non-act-rational equilibrium.

We conclude the paper with a discussion of a few extensions and implications of our model (Section 8): (1) we discuss the empirical predictions of our model; (2) we extend the model to a setup of partial observability; (3) we use our model to study the value of information (see the related literature in the following section) and its relations to Stackelberg stability; and (4) we present ideas for future research.

---

[5] A similar result was obtained in the framework of value of information in Bassan et al. (2003).

2   Related Literature

This paper is related to and inspired by three strands of literature (in addition to the literature on commitments mentioned earlier). The first strand presents an evolutionary foundation for seemingly irrational behavior in strategic environments. Pioneered by Guth & Yaari (1992), the literature on the evolution of preferences shows that if the preferences of an agent are observed by an opponent, then his subjective preferences (which the agent maximizes) may differ from his material preferences (which determine his fitness); see, e.g., Dekel et al. (2007) and Heifetz et al. (2007); Winter et al. (2010) interpret these subjective preferences to be "rational emotions."[6] This literature describes stable outcomes of an evolutionary process in which subjective preferences evolve slowly, and behavior (given the subjective preferences) evolves at a faster pace. Recently, a few papers have used a similar methodology to study the evolution of cognitive biases, such as the "endowment effect" (Frenkel et al., 2014) and bounded forward-looking (Heller, 2015). In this paper we adopt a similar approach, and study evolutionary foundations for ignoring the differences between different games, and bundling them together in a single equivalence set.

The second strand of literature studies *categorization.* It is commonly accepted in the psychology literature that people categorize the world around them in order to help the cognitively limited categorizer to deal with the huge amount of information he obtains from his complex environment. Samuelson (2001) presents a model of cognitively limited agents who bundle together different kinds of bargaining games in order to free up cognitive resources for competitive tournament games. Azrieli (2009) studies games with many players, and shows that categorizing the opponents into a few groups can lead to efficient outcomes. Mohlin (2011) studies the optimal categorization in prediction problems, and Ellison & Holden (2013) study optimal development of categorizing rules in complex environments. Finally, Mengel (2012) studies a learning process in which agents jointly learn which games to bundle together, and how to play in each partition element. The main novelty of this paper relative to this literature is the study of the strategic "commitment" advantages of observable categorization in an interactive environment that includes several possible games.

Jehiel (2005) and Jehiel & Koessler (2008) model players who bundle together different nodes in an extensive-form game into analogy classes, and best-reply to the average opponent's behavior in each class. The key difference between their model and ours is that in theirs the partition of each player is exogenously given, while in our model the partitions are endogenously determined as part of the solution concept. Jehiel (2005, Section 6) writes that "understanding how players categorize contingencies into analogy classes is a very challenging task left for future research." To fill the important gap of understanding how players categorize contingencies into analogy classes, the current paper presents a model in which the categorization is the result of an evolutionary process.

Another novelty of our paper with respect to the literature above is that we introduce an equilibrium concept for an environment with several games, rather than a single game. In many setups, it seems plausible that players' reasoning in one game is influenced by how they played in other games in the past or how they might play in such games in the future. We therefore believe that the concept of an "environment" is a natural concept for studying bounded rationality and it might be useful for other bounded rationality approaches and not merely for rule rationality.

A few recent papers experimentally study categorization of games. Specifically, Grimm & Mengel (2012) study an experiment in which the interactions are drawn randomly from a small set of games in

---

[6]Also related are the papers that study environments in which a player can choose a delegate (with different incentives) to play on his behalf (see, e.g., Fershtman et al. (1991)).

each of 100 rounds. They show that if the environment is not too simple, then there is clear evidence of learning spillovers between the different games. Moreover, Grimm & Mengel (2012) demonstrate that their results "can be captured by a model of categorization where some agents rely on coarse partitions and make choices using a best response correspondence based on the "average game" in each category," and that these partitions depend on the labels of the actions. Both properties hold in the model presented in this paper. Related experimental results can also be found in Huck et al. (2011) and Bednar et al. (2012).

The third strand deals with the *value of information*. In a seminal paper, Blackwell (1953) shows that for single-agent decision problems, more information for the agent is always better. This result can also be extended to constant-sum games, but it is not true, in general, for non-constant-sum games. A classical example of this is given in Akerlof (1970), which shows that providing the seller with private information might render any trade impossible, and thereby reduce the welfare of both the seller and the buyer; see additional examples in Kamien et al. (1990). A few papers characterize specific environments in which information has positive value. Bassan et al. (2003) show that if all games have common interests, then information has a (Pareto-)positive value: more information yields a Pareto-dominant equilibrium (see also Lehrer et al. (2010)). Gossner (2010) shows that having more information is formally equivalent to having a larger set of actions. In Section 8 we show that our model can be reinterpreted to the above framework, and this yields interesting insights about the value of information.

3   Model

Throughout the paper we focus on two-player games and we use index i to denote one of the two players and j to denote his opponent. An environment is a finite set of possible two-player games and a probability function according to which nature chooses the realized game. Formally:

Definition 1. An *environment* $((\Omega, p), A, u)$ is a tuple where:

- (Ω, p) is a finite probability space. That is, Ω is a finite set of states of nature, and p is a probability function over Ω. Without loss of generality, we assume that $p(\omega) > 0$ for each $\omega \in \Omega$.

- $A = A^1 \times A^2$, and each A^i is a finite set of pure actions for player i that are assumed to be available to him at all states of nature.[7]

- $u = \Omega \times A \to \mathbb{R}^2$ is the payoff function that determines the payoff for each player at each state of the world for each action profile.

We assume that both players evaluate random outcomes by using the von Neumann–Morgenstern expected utility extension of u, and, with a slight abuse of notation, we denote by u also this linear extension. Observe that each state of nature ω induces a normal-form game $(A, u(\omega))$ with a set of actions A and a payoff function u.

Remark 1. For simplicity, we only define environments with two players. However, all the definitions and all the results can be extended to n-player environments in a straightforward way.

A *partition* over a set X is a set of non-empty subsets (*partition elements*) of X such that every element in X is in exactly one of these subsets. Let Π_Ω denote the set of all partitions over Ω. Let

---

[7]The assumption that a player has the same set of actions in all states is without loss of generality. If, originally, player i has different actions in different states, this can be modeled by having A_i large enough, relabeling actions in different states, and having several actions yielding the same outcome in some states.

$\bar{\pi} = \{\{\omega\} \,|\, \omega \in \Omega\}$ denote the *finest partition*, and let $\underline{\pi} = \{\Omega\}$ denote the *coarsest partition*. Given a partition $\pi_i \in \Pi_\Omega$ and a state $\omega \in \Omega$, let $\pi_i(\omega)$ be the partition element of π_i that includes state ω.

Each environment induces a two-player two-stage *meta-game*:

1. At stage 1, each player i chooses a partition π_i.

2. At stage 2, nature chooses the state $\omega \in \Omega$ according to p, and each player i observes the chosen partition profile (π_1, π_2) and his partition element $\pi_i(\omega)$, and chooses an action $a^i \in A^i$.

Remark 2.

1. Our definition requires some discussion of the role of partitions in real-life environment and their observability. Moral standards and individual principles are very often transmitted in interactive situations such as negotiations, when one party either tries to persuade the other to act in a certain manner or to threaten the other party against acting in a certain manner. This transmission often takes the form of a verbal exchange of messages combined with either social cues that indicate sincereness or evidence that can verify claims. In many cases, these claims can be naturally interpreted as partitions. For example, a seller of a used car who insists on getting the price quoted in the ad by alluding to the fact that he never bargains as a matter of principle makes a buyer aware of his extremely coarse partition. A buyer of such a car may likewise indicate a moral constraint by arguing that she will never get into a deal with someone who does not show even minimal willingness to compromise. She might refer to verifiable evidence about previous incidents in which negotiations broke down for precisely this reason. Both agents in this example indicate that their behavior is independent of how he (she) is eager to sell (buy), which can be interpreted here as the set of states. Indeed such assertions can be deceptive, but a genuine principle or moral standard will always be more credible than a deceptive one. Religions also impose constraints that can be expressed in a form of partitions. An observant Jew will avoid working on the Sabbath regardless of the financial consequences of such avoidance. An observant Muslim will avoid eating during the day during Ramadan regardless of his/her working schedule that day. Principles derived from religion are much more easily transmitted to other agents than those described earlier in the discussion. Special religious clothing and other observable customs are the source of credibility. In reality they are rarely used to deceive.

2. We interpret the partition chosen by a player as a behavioral phenomenon, and not necessarily as a cognitive limitation. If two games belong to the same element of the partition it does not necessarily mean that the player is unaware of the fact that these are two different games. It does, however, say that the player is bounded to use the same mixed action in both games. This could for instance arise from attributing irrelevance to the difference between these games or because the player is driven by some moral concern that leads him to use the same actions in both games. Note that this interpretation of the role of the partition allows us to assume that while a player bundles two games in the same element of the partition, he can still reasonably believe that his opponent will treat these two games differently.

3. We view the observability of the opponent's partition as a reduced-form model to a richer setup where players either (1) communicate about the strategic setup before playing it, and this reveals information about the opponent's partition, or (2) observe each other's past behavior or a trait that is correlated with the partition. The assumption of perfect observability is taken for tractability,

and it is extended to partial observability in Section 8.2. The first interpretation (pre-play communication) is consistent with experimental findings that observing short pre-play communication allows one to predict players' choices better than a random guess (e.g., Israel et al., 2014). This can also explain observability of partitions that are chosen off the equilibrium path (as such deviations can be detected due to the associated deception in the pre-play communication).

A strategy in the meta-game has two components:

1. A (potentially mixed) partition, i.e., a rule that determines which different situations (games) are bundled together and induce the same behavior.

2. A playing rule. Given the pair of choices of partitions by the two players, and the state of the world (the game played), a player chooses a (potentially mixed) action with one constraint: the same mixed action is chosen for every two states that belong to the same component of his chosen partition.

Formally:

Definition 2. A *(behavior) strategy* for player i in the meta-game is a pair $\tau_i = (\mu_i, \sigma_i)$, where:

1. $\mu_i \in \Delta(\Pi_\Omega)$ is a distribution over the set of partitions.

2. $\sigma_i : \Pi_\Omega \times \Pi_\Omega \times \Omega \to \Delta(A^i)$ is a *playing-rule,* i.e., a mapping that is measurable with respect to the partition of player i: $\sigma_i(\pi_i, \pi_j, \omega) = \sigma_i(\pi_i, \pi_j, \omega')$ whenever $\pi_i(\omega) = \pi_i(\omega')$. That is, playing rule σ_i assigns a mixed action as a function of the partition profile (π_1, π_2) and the partition element of player i that includes the state $\pi_i(\omega)$.

Let $C(\mu)$ denote the support of the distribution μ. Let $x_\tau(\omega) \in \Delta(A)$ be the *interim outcome* of τ - the (possibly correlated) action profile that is induced in state ω given that both players follow the strategy profile τ:

$$x_\tau(\omega)(a_1, a_2) = \sum_{\pi_1 \in C(\mu_1), \pi_2 \in C(\mu_2)} \mu_1(\pi_1) \cdot \mu_2(\pi_2) \cdot \sigma_1(\pi_1, \pi_2, \omega)(a_1) \cdot \sigma_2(\pi_1, \pi_2, \omega)(a_2).$$

Finally, let $\bar{x}_\tau \in \Delta(A)$ be the *ex-ante outcome of τ*, i.e., $\forall a \in A \ \bar{x}_\tau(a) = \sum_{\omega \in \Omega} p(\omega) \cdot x_\tau(\omega)(a)$. With some abuse of notation, let π_i also denote the distribution with mass 1 on partition π_i.

A rule-rational equilibrium is a subgame-perfect equilibrium of the meta-game. Formally:

Definition 3. A strategy profile $\tau^* = ((\mu_1^*, \sigma_1^*), (\mu_2^*, \sigma_2^*))$ is a *rule-rational equilibrium* if for each player i:

1. Each playing-rule is a best reply in each subgame: for each partition profile (π_1, π_2), partition element $\pi_i(\omega)$, and action $a \in A^i$:

$$\sum_{\omega' \in \pi_i(\omega)} p(\omega') \cdot \left(u_i\left(\sigma_i^*(\pi_1, \pi_2, \omega'), \sigma_j^*(\pi_1, \pi_2, \omega') \right) - u_i\left(a, \sigma_j^*(\pi_1, \pi_2, \omega) \right) \right) \geq 0.$$

2. Each distribution of partitions is a best reply: for each chosen partition $\pi_i^* \in C(\mu_i^*)$ and for each partition π_i':

$$\sum_{\pi_j \in C(\mu_j^*)} \mu_j^*(\pi_j) \cdot \sum_{\omega \in \Omega} p(\omega) \cdot \Big(u_i\left(\sigma_i^*(\pi_i^*, \pi_j, \omega), \sigma_j^*(\pi_i^*, \pi_j, \omega) \right)$$
$$-u_i\left(\sigma_i^*(\pi_i', \pi_j, \omega), \sigma_j^*(\pi_i', \pi_j, \omega) \right) \Big) \quad \geq \quad 0.$$

Note that existence of rule-rational equilibria in any environment is immediately implied from the well-known existence result of subgame-perfect equilibria in finite games (Selten, 1975).

Remark 3. We do not interpret rule-rational equilibria to be the result of an explicit payoff maximization of "hyper-rational" agents who choose optimal partitions. Rather, we see the optimization in both stages as arising from an evolutionary process of cultural learning in a large population of agents, namely, the evolutionary game theory approach as pioneered in Maynard-Smith & Price (1973). Each agent in the population is endowed with a type that determines both his rule for bundling games together in a partition and his behavior in the second stage (given the partition). Most of the time agents follow the partition and behavior that they have inherited. Every so often, however, an agent ("mutant") may experiment with a new rule of thumb or a new behavior. The frequency of types evolves according to a payoff-monotonic selection dynamic: more successful types become more frequent. A specific payoff-monotonic dynamics that is often studied in evolutionary models is the "replicator dynamics" (Taylor & Jonker, 1978), according to which the relative change in the frequency of each type is proportional to its average fitness. The replicator dynamics has plausible interpretations both in biological processes and cultural imitation processes (see, e.g., Weibull, 1997, Sections 3.1 and 4.4). Bomze (1986) shows that any asymptotically stable state (i.e., a population starting close to the state eventually converges to it) must correspond to a perfect equilibrium of the underlying game. Hart (2002) studies extensive-form games with perfect information and shows that the evolutionarily stable outcomes coincide with the subgame-perfect equilibria. These results motivate us to present a static reduced-form approach for this evolutionary process (rather than specifying the exact dynamic model), and to use a subgame-perfect equilibrium of the meta-game as the solution concept in the model. We think that this approach sets up a good tradeoff between tractability and rigorousness, even though:

1. Subgame perfection may be too strong in some dynamics (which differ from the replicator dynamics) in which asymptotically stable states are not necessarily perfect (see, e.g., Nachbar, 1990 and Gale et al., 1995).

2. Subgame perfection may be too weak because it does not necessarily imply dynamic stability. A stronger notion that implies dynamic stability in a broad set of payoff monotonic dynamics is evolutionarily stable strategy (Maynard-Smith & Price, 1973). However, our environment typically does not admit any evolutionarily stable strategy, due to the existence of "equivalent" strategies that differ only "off the equilibrium path." Selten (1983) presented the notion of limit evolutionary stability (which refines subgame perfection) that can deal with extensive-form games with multiple stages (see also its correction in Heller, 2014). However, this notion is technically complex (especially when being applied to asymmetric interactions), and some environments still may not admit any limit evolutionarily stable state.

For tractability, we use subgame-perfection as the solution concept, and leave for future research a more complicated model that will either directly study specific selection dynamics or will use a more complex solution concept.

Remark 4. A rule-rational equilibrium differs from standard solution concepts of the incomplete information game with an exogenous information structure. However, a rule-rational equilibrium can be interpreted as a solution concept of an incomplete information game in which players can choose (at the ex-ante stage) how much information to obtain (see Section 8.3 for further discussion).

The following definition divides the rule-rational equilibria into two disjoint subsets: (1) act-rational equilibria in which both players only use the finest partition (which implies playing a Nash equilibrium

in all states), and (2) non-act-rational equilibria in which the interim outcome in at least one state is not a Nash equilibrium.

Definition 4. A rule-rational equilibrium $\tau = ((\mu_1, \sigma_1), (\mu_2, \sigma_2))$ is *act-rational* if $\mu_1 = \mu_2 = \bar{\pi}$. A rule-rational equilibrium is *non-act-rational* if there exists a state $\omega \in \Omega$ such that $x_\tau(\omega)$ is not a Nash equilibrium of the game $(A, u(\omega))$.

Observe that being a non-act-rational equilibrium is slightly stronger than not being an act-rational equilibrium, as the partitions used in the latter case need not be the finest as long as the interim outcome is still a Nash equilibrium in every state.

4   Motivating Examples

This section includes five motivating examples. Example 2 shows an environment in which a coarse partition yields a Pareto-dominant outcome in all states (relative to the set of act-rational payoffs). Example 4 shows a similar result for team-effort games. Example 3 demonstrates an environment that does not admit act-rational equilibria. Finally, we demonstrate that non-act-rational equilibria can fit experimentally observed behavior in the "Ultimatum game" (Example 5) and the "Chain store game" (Example 6).

4.1   Non-Act Rationality Pareto-Improves Payoffs

Our first example demonstrates two interesting points regarding non-act-rational equilibria.

1. A non-act-rational equilibrium yields a strict Pareto-improvement over the unique act-rational equilibrium.

2. The partitions in this equilibrium are strictly in between the coarsest and the finest partitions, and they have a natural interpretation (players pay attention only to the total gain from trade).

Example 2. Consider an environment in which each player has two actions: T (interpreted as trade) and N (non-trade). A trade interaction occurs iff both agents agree to trade. Nature chooses with equal probability one of the three states: in state ω_1 (ω_2) trade is very profitable to player 1 (2) but slightly unprofitable to player 2 (1); in state ω_3 trade is unprofitable to both players. The different payoffs in each state are described in Table 2.

Table 2: Trade Game

ω_1	T	N
T	5,-1	0,0
N	0,0	0,0

ω_2	T	N
T	-1,5	0,0
N	0,0	0,0

ω_3	T	N
T	-5,-5	0,0
N	0,0	0,0

The unique act-rational equilibrium of the game is the one in which the players do not trade in all states and earn zero. The situation resembles a double-sided "lemon market" (Akerlof, 1970), and the adverse selection prevents efficient trade. The game admits also a Pareto-dominant non-act-rational equilibrium in which both players pay attention only to the total gain from trade (i.e., to the sum of payoffs), while ignoring the details of how the total gain is divided between the two players (i.e., whether one of them owns a "lemon"). Specifically, in this equilibrium both players choose the partition $\{\{\omega_1, \omega_2\}, \omega_3\}$, and they both choose T if the state is in $\{\omega_1, \omega_2\}$, and choose N in ω_3. If any player

deviates to a different partition, then both players choose N in all states. This equilibrium yields both players an expected payoff of $\frac{4}{3}$.

4.2 An Environment without Act-Rational Equilibria

The following example demonstrates an environment that does not admit any act-rational equilibrium.

Example 3. Let $\Omega = \{\omega_1, \omega_2\}$, and let A and each $u(\omega_i)$ be described in Table 3.

Table 3: Environment without Act-Rational Equilibria

ω_1	L	R
C	9,9	0,7
D_1	10,0	1,1
D_2	7,0	-2,1

ω_2	L	R
C	9,9	0,7
D_1	7,0	-2,1
D_2	10,0	1,1

We now show that for every $\frac{1}{3} < p(\omega_1) < \frac{2}{3}$, the environment admits no act-rational equilibrium. Observe that each of the two games admits a dominant action to player 1 (D_1 in state ω_1 and D_2 in state ω_2), which induces a unique Nash equilibrium ((D_1, R) in ω_1 and (D_2, R) in ω_2) with payoff $(1, 1)$. This implies that the unique payoff of each player in any act-rational equilibrium is one. However, if player 1 deviates to the coarse partition, then action C becomes the dominant action in the induced subgame with incomplete information, and this implies that the unique Nash equilibrium of this subgame is (C, L), which yields both players a payoff of 9. This shows that the environment does not admit any act-rational equilibrium. An example of a non-act-rational equilibrium in this environment is as follows. Let $\pi_1^* = \underline{\pi}$, $\pi_2^* = \overline{\pi}$, and let σ^* be equal to (C, L) either on the equilibrium path or if player 2 deviated to $\underline{\pi}$; otherwise (i.e., if player 1 deviated to $\overline{\pi}$) let σ^* be equal to (D_1, R) in ω_1 and (D_2, R) in ω_2 .

4.3 Rule Rationality and Team-Effort Games

The following example shows that rule rationality allows one to achieve maximal efficiency in a family of team-effort games in which act rationality allows only partial efficiency.

Example 4. Consider a project that includes two tasks. Each of these tasks is assigned to a different player. The project succeeds if and only if both tasks are successful. A player succeeds in a task with certainty if he exerts effort and succeeds with probability α if he does not. Exerting effort has a cost of 0.5. If the project succeeds they each get a payoff of one. The environment consists of all the games arising from some $0 \le \alpha \le 1$. That is, let $\Omega = [0, 1]$, where $\omega \in \Omega$ is interpreted as the value of α, and let p be a probability distribution over Ω.[8]

The environment admits multiple act-rational equilibria. The least efficient equilibrium is the one in which the players never exert effort. This is an equilibrium because not exerting effort is the unique best-reply to a partner who does not exert effort: it yields α^2, while exerting effort yields only $\alpha - 0.5 < \alpha^2$. The most efficient act-rational equilibrium is the one in which each player exerts effort for each $\alpha \le 0.5$. This is an equilibrium because for any $\alpha \le 0.5$, exerting effort is a best-reply to exerting effort: it yields $1 - 0.5 = 0.5$, while not exerting effort yields $\alpha \le 0.5$. The environment does not admit a more efficient act-rational equilibrium because not exerting effort yields a strictly higher payoff regardless of the opponent's action for any $\alpha > 0.5$.

---

[8]In this example we slightly extend the model by having a continuum of states.

Assume that the expected value of α is less than 0.5 ($\mathbf{E}_p(\alpha) \leq 0.5$). Under this assumption, the environment admits a non-act-rational equilibrium that yields the fully efficient outcome in which both agents always exert effort. In this equilibrium, both players chooses the coarsest partition, and then exert effort in the second stage. If any player deviated to a finer partition, then neither player exerts effort in the second stage. This strategy profile is an equilibrium because: (1) deviating at the second stage by not exerting effort yields a payoff of $\mathbf{E}_p(\alpha)$, which is assumed to be at most 0.5 (the equilibrium payoff); and (2) deviating at the first stage (i.e., choosing a finer partition) yields a payoff of $\mathbf{E}_p(\alpha^2) < \mathbf{E}_p(\alpha) \leq 0.5$.

In Section 7.3 we use an extension of the model to study related team-effort games with private information about the cost of exerting effort.

4.4   Environments with Extensive-Form Games

4.4.1   Extending the Model to Extensive-Form Games

To simplify the presentation of the general model in Section 3 we defined an environment consisting of normal-form games. The following two examples differ somewhat in that in each state of nature players interact in sequential games. In what follows we briefly describe how to adapt the model in order to deal with extensive-form games. For brevity we do not present all the formal details, as the extension is relatively straightforward.

We redefine an *environment* to be a tuple $((\Omega, p), \Gamma, u)$ where: (Ω, p) is a finite probability space, as before; Γ is a "game structure", i.e., an extensive-form game without the component that defines the utility of each player in each terminal node; and u is a mapping that assigns a payoff to each player at each state of the world for each terminal node of Γ.

As before, we require a rule-rational equilibrium to be a subgame-perfect equilibrium of the meta-game. Finally, we slightly relax the definition of a non-act-rational equilibrium to require that there be a state in which the induced behavior in the extensive-form game is not a subgame-perfect equilibrium.

4.4.2   Ultimatum Game and Sequential Bargaining

The following example describes a situation in which players play a sequential bargaining game, and the length of the game depends on the state of nature.[9]

Example 5. Consider an environment that includes two states with equal probability: $\Omega = \{\omega_1, \omega_2\}$. In each state the two risk-neutral players have to agree how to divide one dollar between them, and they get zero if they do not reach an agreement. In state ω_1 they play the "ultimatum" game: player 1 offers a division, which player 2 either accepts or rejects, and the game ends. In state ω_2 they play sequential bargaining with two rounds: if player 2 rejects player 1's offer, then player 2 makes a counteroffer in round 2, which player 1 either accepts or rejects, and then the game ends. [10]

One can see that the environment admits the following two pure rule-rational equilibria:

1. Act-rational equilibrium. Both players choose the fine partition. Player 1 offers zero to player 2 in ω_1 and offers one in ω_2. Player 2 accepts these offers. Off the equilibrium path, player 2 (with a fine partition) accepts any offer in ω_1 and rejects any offer below 1 in ω_2. If a second round is played in ω_2, then player 2 offers 0 to player 1, and player 1 accepts any offer. If player 1 chooses the coarse partition, then player 1 offers zero to player 2 in both states and player 2 plays the same

---

[9]In this example we slightly extend the model by having a continuum of pure actions for each player.

[10]As stated above, the players play a different game-structure in each state. However, this can also be modeled with the same two-round game structure, where in ω_1 the choices in round 2 do not influence the payoffs.

as before (accepts the offer in ω_1 and rejects it in ω_2). If player 2 chooses the coarse partition, then player 1 offers 0.5 in both states, and player 2 accepts any offer of at least 0.5 (and the players behave as before if a second round is reached).

2. Non-act-rational equilibrium. Both players choose the coarse partition. Player 1 offers player 2 a payoff of 0.5 in both states. Player 2 accepts this offer. Off the equilibrium path, player 2 rejects an initial offer iff it is below 0.5. If a second round is played, then player 2 offers zero, and player 1 accepts it. If player 1 chooses the fine partition, then the players still play the same (player 1 offers 0.5 in both states, and player 2 accepts any offer of at least 0.5). If player 2 chooses the fine partition, then player 1 offers 0 to player 2 in both states, and player 2 accepts any offer in ω_1 and rejects any offer in ω_2 (and they play the same as before in the second round). Note that the outcome in state ω_2 is not a Nash equilibrium (the best reply of player 2 in state ω_2 is to reject the offer of 0.5), and the outcome in state ω_1 is an imperfect Nash equilibrium.

Note that both equilibria yield both players the same ex-ante payoff (0.5 for each player).[11] The coarse partition in the non-act-rational equilibrium may be interpreted as the players not paying attention to the exact length of the bargaining, and behaving the same in all bargaining situations. Note that the observable behavior (on the equilibrium path) of the non-act-rational equilibrium fits the following stylized facts about the experimental behavior in the "ultimatum" game (see Roth et al., 1991 and the survey of Güth & Kocher, 2014): (1) subjects usually offer 40–50% of the pie in the first round, (2) such offers are almost always accepted, and (3) acceptance rates decrease with smaller offers. [12]

4.4.3 Chain-Store Game

The following example deals with the chain-store game (Selten, 1978): a monopolist faces a finite sequence of potential entrants, and if any of them chooses to enter, the monopolist may "fight" the entrant by using predatory pricing. The example shows a non-act-rational equilibrium in which the monopolist chooses not to pay attention to the exact number of remaining potential entrants, and fights all entrants. This behavior roughly fits the observed behavior in experiments (subjects playing the monopolist tend to fight entrants in most rounds of the interaction; see Jung et al., 1994, and qualitatively similar behavior in a related repeated "trust" interaction in Neral & Ochs, 1992.) Observe that in any act-rational equilibrium the monopolist never fights any entrant.[13]

Example 6. Consider an environment with a countable number of states: $\Omega = \mathbb{N}$. Let p be any distribution satisfying: (1) full support - $p(\omega) > 0 \ \forall \omega \in \mathbb{N}$, and (2) "continuation" probability that is high enough - $\forall k \in \mathbb{N}, \ p(\omega = k | \omega \geq k) < 2/3$. In particular these requirements are satisfied by any geometric distribution with high enough expectation. The state ω describes the number of potential competitors that the *monopolist* (player 1) faces. The game lasts ω rounds. To simplify notation and fit our two-player setup, we represent all potential competitors by a single "impatient" player 2, the *competitor*. At each round, player 2 chooses In or Out. Out yields a stage payoff of 5 to the monopolist and 1 to the competitor. If he chose In, then the monopolist chooses Fight (predatory pricing) or Adapt.

---

[11] If the players have modest risk-aversion or a discount factor slightly less than one, then the second equilibrium remains very similar, except that the offer in the first round is slightly less than 0.5.

[12] Unlike the simple behavior in the above non-act-rational equilibrium, lab participants playing the alternating offer game do respond to the game horizon (see, for example, Ochs & Roth, 1989). However, this dependence is much smaller than the prediction of the subgame-perfect equilibrium and the agreed-upon allocations are strongly tilted towards the 50:50 share. Our example explains how such drift towards the 50:50 share can arise from a commitment to ignore the horizon. One could construct a more complex example in which a non-act-rational equilibrium combines a tendency towards the 50:50 share with some response to the game horizon, which is closer to the observed experimental results.

[13] In this example we slightly extend the model by allowing the set of states to be countable (rather than finite).

Fight yields stage payoff of 0 to both players, and Adapt yields 2 to both players. The total game payoff of the monopolist is the discounted sum of stage payoffs for some high enough discount factor $0 << \lambda \leq 1$. The total payoff of the competitor is a discounted sum of the payoffs with a low discount factor $0 \leq \lambda << 1$. In the classical setup in which the monopolist knows the state of nature, the game admits a unique subgame-perfect equilibrium: the competitor always plays In, and the monopolist always plays Adapt. In our setup, however, there is an additional non-act-rational equilibrium that yields a better payoff to the monopolist. In this equilibrium the monopolist chooses the coarsest partition (i.e., he does not pay attention to the exact length of the game), and plays Fight following any entry. The competitor chooses the finest partition and plays Out, unless he observes an out-of-equilibrium history in which the monopolist played Adapt (and in this case he plays In).

5   Results

This section characterizes rule-rational equilibria in families of environments that satisfy certain properties. In Section 5.1 we show that rule-rational equilibria are closely related to correlated equilibria in environments that include best-response equivalent games. In Section 5.2 we show that environments that only include constant-sum games essentially admit only act-rational equilibria (but this is not true for the larger family of strategically zero-sum games). Section 5.3 shows that environments with common interests admit a Pareto-dominant act-rational equilibrium. Finally, in Section 5.4 we study environments that assign high enough probability to a game with an undominated, strict equilibrium.

5.1   Environments with a Uniform Best-Response Correspondence

Two games with the same set of actions, (A, u) and $(A, \tilde{u})$, are *best-response equivalent* if they have the same best-response correspondence. Formally, for each player $i \in \{1, 2\}$ and each mixed action $x_{-i} \in \Delta(A_{-i})$,

$$\mathrm{argmax}_{a_i \in A_i} u(a_i, x_{-i}) = \mathrm{argmax}_{a_i \in A_i} \tilde{u}(a_i, x_{-i})$$

(see Morris & Ui, 2004 and the references there). It is well known that best-response equivalent games admit the same sets of Nash equilibria and correlated equilibria (Morris & Ui, 2004). We say that an environment has a *uniform best-response correspondence* if it includes only best-response equivalent games.

Recall that a correlated equilibrium is a Nash equilibrium of an extended game that includes a pre-play stage in which each player observes a private signal (and the signals may be correlated). Formally:

Definition 5. (Aumann, 1974, as reformulated in Osborne & Rubinstein, 1994) A correlated equilibrium of a two-player normal-form game $G = (A, u)$ is a tuple $\zeta = \left((\tilde{\Omega}, \tilde{p}), (\tilde{\pi}_1, \tilde{\pi}_2), (\tilde{\eta}_1, \tilde{\eta}_2) \right)$ that consists of:

- A finite probability space $(\tilde{\Omega}, \tilde{p})$,

- A partition profile $(\tilde{\pi}_1, \tilde{\pi}_2) \in \Pi_{\tilde{\Omega}} \times \Pi_{\tilde{\Omega}}$,

- A strategy profile $(\tilde{\eta}_1, \tilde{\eta}_2)$, where each $\tilde{\eta}_i : \Omega \to A_i$ is a $\tilde{\pi}_i$-measurable strategy,

such that for every player i and every π_i-measurable strategy $\eta_i' : \Omega \to A_i$, we have

$$\sum_{\omega \in \tilde{\Omega}} \tilde{p}(\omega) \cdot u_i(\tilde{\eta}_i(\omega), \tilde{\eta}_j(\omega)) \geq \sum_{\omega \in \tilde{\Omega}} \tilde{p}(\omega) \cdot u_i(\eta_i'(\omega), \tilde{\eta}_j(\omega)).$$

We say that an environment is *redundant* if it includes only exact copies of the same game (i.e., $u(\omega) = u(\omega')$ for each $\omega, \omega' \in \Omega$). It is immediate that any redundant environment has a uniform best-response correspondence. The following example demonstrates that a redundant environment may admit a rule-rational equilibrium with a payoff outside the convex hull of the set of Nash equilibrium payoffs.

Example 7. Consider the following environment: $\Omega = \{\omega_1, \omega_2, \omega_3\}$, p is uniform, and (A, u) is the same "Chicken" (or "Hawk-Dove") game in all states (see Table 4). Consider the following rule-rational equilib-

Table 4: Payoffs of the "Chicken" Game in All States

	C	D
C	3,3	1,4
D	4,1	0,0

rium: player 1 chooses partition $\pi_1^* = \{\{\omega_1, \omega_2, \}\omega_3\}$ and player 2 chooses partition $\pi_2^* = \{\omega_1, \{\omega_2, \omega_3\}\}$. On the equilibrium path each player plays C when he does not know the exact state of nature, and plays D otherwise. Off the equilibrium path, if a player chooses any other partition, the deviator plays C and his opponent plays D.[14] This rule-rational equilibrium (which "imitates" a correlated equilibrium of the "Chicken" game) yields each player a payoff of $8/3$ (and this is outside the convex hull of the set of Nash equilibria).

The following proposition shows that an environment with a uniform best-response correspondence can only induce behavior that is consistent with a correlated equilibrium. Formally:

Proposition 1. *Let* $((\Omega, p), A, u)$ *be an environment with a uniform best-response correspondence. Then:*

1. *Any Nash equilibrium can be induced by an act-rational equilibrium.*

2. *The induced aggregate action profile of every rule-rational equilibrium is always a correlated equilibrium (of every game in the environment).*

Proof. □

1. Let $x \in \Delta(A_1) \times \Delta(A_2)$ be a Nash equilibrium of the games $(A, u(\omega))_{\omega \in \Omega}$. It is immediate that any arbitrary choice of partitions, followed by playing x in all states (regardless of the chosen partitions), yields an act-rational equilibrium.

2. Let $\tau^* = ((\mu_1^*, \sigma_1^*), (\mu_2^*, \sigma_2^*))$ be a rule-rational equilibrium. Intuitively, a mediator can "mimic" the choice of the state of nature and of the partition that each player uses, and inform each player about his partition and about the partition element that includes the state. Formally, we construct a correlated equilibrium $\zeta_{\tau^*} = ((\tilde{\Omega}, \tilde{p}), (\tilde{\pi}_1, \tilde{\pi}_2), (\tilde{\eta}_1, \tilde{\eta}_2))$ (of every game in the environment) as follows. Let $\tilde{\Omega} = \Omega \times \Pi_\Omega \times \Pi_\Omega$, and let $\tilde{p}(\omega, \pi_1, \pi_2) = p(\omega) \cdot \mu_1^*(\pi_1) \cdot \mu_2^*(\pi_2)$. For each player i, let $\tilde{\pi}_i((\omega_0, \pi_1, \pi_2)) = \{(\omega, \pi_1, \pi_2) | \omega \in \pi_i(\omega_0)\}$, and let $\tilde{\eta}_i(\omega_0) = \sigma_i^*(\pi_1, \pi_2, \omega_0)$. The fact that τ^* is a rule-rational equilibrium immediately implies that ζ_{τ^*} is a correlated equilibrium.

The following proposition shows that the converse is also true: any correlated equilibrium (which is not dominated by the set of Nash equilibria) can be induced by a redundant environment.

---

[14]For brevity, throughout the proofs in this paper we specify only the strategies after a single player deviates. The play following simultaneous deviations of both players at the partition stage can be determined arbitrarily as any Nash equilibrium of the induced subgame.

Definition 6. A correlated equilibrium is undominated if for each player i there exists a Nash equilibrium that yields player i a weakly worse payoff.

Proposition 2. *For every undominated correlated equilibrium ζ in game (A, u_0), there exists a redundant environment $((\Omega, p), A, u \equiv u_0)$ and a rule-rational equilibrium that induces ζ as the aggregate outcome.*

Proof. Let $\zeta = \left((\tilde{\Omega}, \tilde{p}), (\tilde{\pi}_1, \tilde{\pi}_2), (\tilde{\eta}_1, \tilde{\eta}_2) \right)$ be a correlated equilibrium of the game (A, u_0). Consider the following strategy profile $\tau_\zeta = ((\mu_1, \sigma_1), (\mu_2, \sigma_2))$ in the redundant environment $((\Omega, p), A, u \equiv u_0)$: each μ_i assigns mass 1 to $\tilde{\pi}_i$, and $\sigma_i = \tilde{\eta}_i$ if the realized partition profile is $(\tilde{\pi}_1, \tilde{\pi}_2)$. If player i deviates and plays a different partition, both players play in all states the Nash equilibrium that yields player i a worse payoff. It is immediate to see that τ_ζ is a rule-rational equilibrium. $\qquad\square$

Remark 5. Dominated correlated equilibria (i.e., correlated equilibria that are not undominated) cannot be implemented by a rule-rational equilibrium of redundant environments. The sketch of the proof is as follows. Assume to the contrary that a dominated correlated equilibrium ζ is implemented as a rule-rational equilibrium of an environment that includes only copies of the same game. Let i be a player that obtains a strictly higher payoff in all Nash equilibria than in ζ. Consider a deviation of player i in which he chooses the finest partition. Consider the subgame following this deviation and an arbitrary partition element of the opponent $\pi_j(\omega)$. Observe that the average mixed action played by player i conditional on the state being $\pi_j(\omega)$ must be a best reply to the mixed action played by his opponent in $\pi_j(\omega)$, and vice versa. This implies that the payoff of player i after this deviation is in the convex hull of his Nash equilibrium payoffs, and thus it is strictly higher than his payoff in ζ. An example of a dominated correlated equilibrium can be found in Evangelista & Raghavan (1996, page 40).

Finally in Appendix A we show that the implementation of correlated equilibria does not depend on using "tailored" probability distributions. Specifically, we show that any undominated correlated equilibrium can be approximately induced by any large enough "atomless" environment (an environment in which each state has a small probability).

The results of this section show that rule rationality induces the play of correlated equilibria in environments of strategically equivalent situations (which may be sub-environment embedded in a larger class of potentially very different games). Correlated equilibrium is an important solution concept in game theory, but its implementation is demanding: it requires players to rely on a reliable mediator for doing a joint lottery, and giving each player different partial information about the lottery's result. In many strategic situations, reliable mediators are not available, and it is not clear how the correlation can be induced.[15] Our results present a novel explanation: correlated equilibria can arise as a result of rule of thumbs that players use to bundle together similar situations.

5.2   Constant-Sum and Strategically Zero-Sum Environments

Environment $((\Omega, p), A, u)$ is constant-sum if for each state $\omega \in \Omega$, the game $(A, u(\omega))$ is a constant-sum game. In this section we show that constant-sum environments essentially admit only act-rational equilibria. Specifically, we show that a non-act-rational equilibrium is not robust to an arbitrarily small probability that the opponent may overlook a deviation at the partition stage.

---

[15] Alternative models that induce correlation without mediators are Hart & Mas-Colell (2000), in which it is induced by a long history of play, and Ben-Porath (2003) (see related models in the references within), in which the correlation is induced by a complicated cheap-talk protocol when there are at least three players. See also Aumann (1987)'s interpretation of correlated equilibria as an expression of Bayesian rationality.

We begin by defining the notion of robustness of an equilibrium. The notion assumes that when a player deviates while choosing a partition, then: (1) with very high probability the deviation is observed by the opponent and the players follow the off-equilibrium-path strategies, and (2) with very low probability the deviation is unobserved by the opponent, and the deviating player (who is aware of the fact that his deviation was unobserved) can play any playing rule that is consistent with his partition. Formally:

Definition 7. Rule-rational equilibrium $\tau^* = ((\mu_1^*, \sigma_1^*), (\mu_2^*, \sigma_2^*))$ is *robust* if there exists $\bar{\epsilon} > 0$ such that for each $\epsilon < \bar{\epsilon}$, for each player i, for each chosen partition $\pi_i^* \in C(\mu_i^*)$, for each partition π_i' and for each mapping $\beta_i' : \Pi_\Omega \times \Omega \to \Delta(A^i)$ that is measurable with respect to π_i' (i.e., $\sigma_i'(\pi, \omega) = \sigma_i'(\pi, \omega')$ whenever $\pi_i'(\omega) = \pi_i'(\omega')$):

$$\sum_{\pi_j^* \in C(\mu_j^*)} \mu_j^*(\pi_j^*) \cdot \sum_{\omega \in \Omega} p(\omega) \cdot u_i\left(\sigma_i^*\left(\pi_i^*, \pi_j^*, \omega\right), \sigma_j^*\left(\pi_i^*, \pi_j^*, \omega\right)\right) \geq$$

$$\sum_{\pi_j^* \in C(\mu_j^*)} \mu_j^*(\pi_j^*) \cdot \left((1-\epsilon) \cdot \sum_{\omega \in \Omega} p(\omega) \cdot u_i\left(\sigma_i^*\left(\pi_i', \pi_j^*, \omega\right), \sigma_j^*\left(\pi_i', \pi_j^*, \omega\right)\right) \right.$$

$$\left. + \epsilon \cdot \sum_{\omega \in \Omega} p(\omega) \cdot u_i\left(\beta_i'\left(\pi_j^*, \omega\right), \sigma_j^*\left(\pi_i^*, \pi_j^*, \omega\right)\right) \right).$$

For each state $\omega \in \Omega$, let $v(\omega)$ be the interim value of the game $(A, u(\omega))$ (i.e., the unique payoff profile that is induced by all Nash equilibria of $(A, u(\omega))$), and let $v(\Omega) = \sum p(\omega) \cdot v(\omega)$ be the *ex-ante* value of the environment. The following result shows that non-act-rational equilibria are not robust in constant-sum environments.

Proposition 3. *Let $((\Omega, p), A, u)$ be a constant-sum environment. Then:*

1. *The environment admits an act-rational equilibrium.*

2. *All rule-rational equilibria yield the same ex-ante payoff, $v(\Omega)$.*

3. *Any non-act-rational equilibrium is not robust.*

Proof.

1. Consider the strategy in which each player chooses the finest partition and then plays a minimax strategy in each game (regardless of the partition his opponent chose). It it immediate that this strategy profile is an act-rational equilibrium.

2. This is immediately implied by the fact each player may deviate by choosing the finest partition, and playing a minimax strategy in each game, which yields an *ex-ante* payoff of at least $v(\Omega)$.

3. For every $\epsilon > 0$, a player can strictly earn by deviating from a rule-rational equilibrium into choosing the finest partition, and playing a minimax strategy in each game. The deviator will earn the same *ex-ante* payoff if the deviation is observed, and will earn a strictly higher payoff otherwise (as there is a state in which with positive probability the players do not play a Nash equilibrium, and one of the player can earn by best-replying to its opponent's play).

$\square$

Next we demonstrate that constant-sum environments may admit a non-robust non-act-rational equilibrium.

Example 8. Let $\Omega = \{\omega_1, \omega_2\}$, $p(\omega_1) = 0.5$, and let A and each $u(\omega)$ be defined according to Table 5. Consider the following strategy profile. Both players choose the coarsest partition ($\forall i \ \pi_i^* = \underline{\pi}$) and they

Table 5: Payoffs of a Zero-sum Environment

	ω_1				ω_2	
	L	R			L	R
T	1,-1	0,0		T	-1,1	0,0
D	0,0	0,0		D	0,0	0,0

both play (T, L) on the equilibrium path, and (D, R) off the equilibrium path (if any player deviated and chose a different partition). It is immediate to see that this profile is a non-act-rational equilibrium. Note that this equilibrium yields the players in each state ω_i a payoff different from the interim value $v(\omega_i) = 0$, while the ex-ante value, $v(\Omega) = 0$, is unchanged.

Two games are von Neumann–Morgenstern (VNM) equivalent if, for each player, the payoff function in one game is equal to a constant times the payoff function in the other game, plus a function that depends only on the opponents' strategies. Note that VNM equivalence implies best-response equivalence, but the converse is not necessarily true (Morris & Ui, 2004). A game is *zero-sum equivalent* if it is VNM equivalent to a zero-sum game.[16] An environment $((\Omega, p), A, u)$ is zero-sum equivalent if for each state $\omega \in \Omega$, the game $(A, u(\omega))$ is zero-sum equivalent. Note that each constant-sum environment is zero-sum equivalent but the converse is not necessarily true. The following example demonstrates that Proposition 3 cannot be extended to zero-sum equivalent environments. Specifically, it shows a zero-sum equivalent environment that admits a robust non-act-rational equilibrium, and does not admit any act-rational equilibrium. The basic intuition is the (well-known) fact that replacing one of the matrices of a Bayesian game with a VNM-equivalent matrix may change the set of equilibria.

Example 9. Let $\Omega = \{\omega_1, \omega_2\}$, $p(\omega_1) = 0.5$, and let A and each $u(\omega)$ be defined as in Table 6. First,

Table 6: Payoffs of a Strategically-Zero-Sum Environment

	ω_1				ω_2	
	L	R			L	R
T	0,0	0,0		T	11,-1	-1,1
D	3,-3	4,-4		D	9,1	1,-1

note that each game is zero-sum equivalent: $(A, u(\omega_1))$ is a zero-sum game, and $(A, u(\omega_2))$ is derived from a zero-sum "matching pennies" game by adding 10 to the first column of the matrix payoff of player 1 (Moulin & Vial, 1978 show that this implies that $(A, u(\omega_2))$ is strategically zero-sum). Next, note that each game admits a unique Nash equilibrium: (D, L) in ω_1 and $((0.5, 0.5), (0.5, 0.5))$ in ω_2. These action profiles yield player 1 an ex-ante payoff of $4 = 0.5 * 3 + 0.5 * 5$. We first show that the environment does not admit any act-rational equilibrium. Assume to the contrary that these action profiles were the interim outcome of an act-rational equilibrium. Note that player 1 can deviate by choosing the coarsest partition. Given this deviation, playing D is a strictly dominant action for player 1, and this implies that playing (D, L) is the unique Nash equilibrium of the induced game, and that it yields player 1 an ex-ante payoff of 6. This contradicts the assumption that an act-rational equilibrium exists.

---

[16]Our notion of zero-sum equivalent games slightly refines Moulin & Vial's (1978) notion of strategically zero-sum games, which relies on a better-reply equivalence (Morris & Ui, 2004).

Next, consider the strategy profile in which player 1 chooses the coarse partition and plays D, while player 2 chooses the fine partition and plays L. This profile yields an ex-ante payoff profile of $(6, 1)$. If player 1 deviates to the fine partition, then the players play the unique Nash equilibrium in each state (which yields player 1 only 4). Player 1 plays D (which is a dominant action given his coarse partition) also if his opponent chooses the fine partition. It is immediate that this strategy profile is a robust non-act-rational equilibrium.

5.3 Environments with Common Interests

A game (A, u_0) has *common interests* (Aumann & Sorin, 1989) if there is an action profile $\tilde{a} \in A$ that Pareto-dominates all other action profiles: $\forall i \in \{1, 2\}$, $a \in A$, $u_i(\tilde{a}) > u_i(a)$. An environment $((\Omega, p), A, u)$ has common interests if it includes only games with common interests. The following proposition shows that rule rationality is very limited in environments with common interests, in the sense that any non-act-rational equilibrium is strictly Pareto-dominated by an act-rational equilibrium.[17] The proposition is related to an existing result of Bassan et al. (2003), which shows that information has a (Pareto) positive value if each game has common interests.

Proposition 4. *Let $((\Omega, p), A, u)$ be an environment with common interests. Then there exists an act-rational equilibrium that strictly Pareto-dominates all non-act-rational equilibria.*

Proof. It is immediate that the strategy profile in which players choose the finest partition and then play the Pareto-optimal action profile in every game is an act-rational equilibrium. Consider any non-act-rational equilibrium. This equilibrium induces in some $\omega_0 \in \Omega$ an outcome that is not a Nash equilibrium. The fact that each game has common interests implies that the interim payoff of the non-act-rational equilibrium is weakly dominated by the interim payoff of the above act-rational equilibrium, and that it is strictly dominated in ω_0. $\square$

We conclude this section by demonstrating the existence of Pareto-dominated non-act-rational equilibria in environments with common interests.

Example 10. Let $\Omega = \{\omega_1, \omega_2\}$, and let A and each $u(\omega_i)$ be defined as in Table 7. One can see that the

Table 7: Payoffs of an Environment with Common Interests

	ω_1				ω_2	
	L	R			L	R
T	5,5	0,0		T	5,5	0,0
M	6,6	0,0		M	0,0	0,0
D	0,0	1,1		D	0,0	1,1

following strategy profile is a non-act-rational equilibrium for every $p(\omega_1) < 5/6$: the players choose the coarsest partition ($\pi_i = \underline{\pi}$) and play (T, L) on the equilibrium path, and (D, R) if any player deviated by choosing a different partition. Note that this equilibrium is Pareto-dominated by the act-rational equilibrium in which both players choose the finest partition, and play (M, L) in ω_1 and (T, L) in ω_2.

---

[17]In addition, one can show that all the non-act-rational equilibria in common-interest environments do not satisfy the forward-induction refinement (Govindan & Wilson (2008)), while the Pareto-dominant act-rational equilibrium satisfies it.

5.4 Environments with Global Rules

In this section we show that in environments of sufficiently similar games a strict equilibrium in one of these games can induce a global rule in the sense that the same profile will be played regardless of the state of the world.

A *global-rule equilibrium* is a rule-rational equilibrium in which both players use the coarsest partition. In what follows we consider an environment that includes a game g^* with a strict equilibrium a that dominates other strict equilibria. The first result shows that if the other games in the environment are sufficiently similar to g^*, and if the probability of g^* is large enough, then the environment admits a global-rule equilibrium in which players play action a in all states. Note that if there is a state in which a is not a Nash equilibrium, then this is a non-act-rational equilibrium. Formally:

Definition 8. Let $\delta > 0$. A pure action profile a in a game (A, u) is a δ-*strict Nash equilibrium* if for each player i and each action $b^i \in A^i$, $u^i(a) \geq u^i(b^i, a^j) + \delta$.

Proposition 5. *Let $\epsilon, \delta > 0$. Let (A, u_0) be a game with δ-strict equilibria (a, b_1, b_2) such that $\forall i \in \{1, 2\}$, $u_0^i(a) > u_0^i(b_i) + \delta$. Let $((\Omega, p), A, u)$ be an environment such that: (1) it includes game (A, u_0): $\exists \omega_0 \in \Omega$ such that $u(\omega_0) = u_0$; (2) the payoff matrices in all states are ϵ-close to u_o: $\forall \omega \in \Omega$, $a \in A$ $i \in \{1, 2\}$: $\left| u^i(\omega)(a) - u_0^i(a) \right| \leq \epsilon$; and (3) $p(\omega_0) \geq \frac{\epsilon - \delta}{\epsilon}$. Then the environment admits a rule-rational equilibrium in which players choose the coarsest partition and play a in all states.*

Proof. Consider the following strategy profile. The players choose the coarsest partition and play a. If player i deviates and chooses a different partition, then the players play profile b_i in the partition element that includes ω_0, and they play an arbitrary equilibrium in the subgame that is induced by the choice of partitions and by fixing the play of profile b_i in ω_0. Note that any unilateral deviation in the partition element that includes ω_0 yields the deviator a loss of at least δ utility points in ω_0 and a profit of at most $\epsilon - \delta$ in any other state. Thus, this deviation yields an expected profit of at most

$$\left(1 - \frac{\epsilon - \delta}{\epsilon}\right) \cdot (\epsilon - \delta) - \frac{\epsilon - \delta}{\delta} \cdot \delta = \frac{\delta}{\epsilon} \cdot \epsilon - \delta = 0.$$

This implies that the above strategy is an equilibrium of the subgame that is induced by the chosen partitions also without fixing profile b_i in ω_0. Observe that a player who deviates at the partition stage loses at least δ in ω_0 and gains at most $\epsilon - \delta$ in any other state; thus the maximal profit from a deviation from the strategy profile is

$$\left(1 - \frac{\epsilon - \delta}{\epsilon}\right) \cdot (\epsilon - \delta) - \frac{\epsilon - \delta}{\epsilon} \cdot \delta = 0.$$

This implies that the above strategy profile is a rule-rational equilibrium. $\square$

The following corollary shows that if g^* (a game with a strict equilibrium a that dominates other strict equilibria) has high enough probability, then the environment admits a rule-rational equilibrium in which players bundle all the other "unlikely" games with the prominent game g^*, and play a in all games. Formally:

Corollary 1. *Let (A, u_0) be a game with strict equilibria (a, b_1, b_2) such that $\forall i \in \{1, 2\}$, $u_i(a) > u_i(b_i)$. Then there exists $0 < p_0 < 1$ such that each environment $((\Omega, p), A, u)$ that satisfies $\omega_0 \in \Omega$, $u(\omega_0) = u_0$, and $p(\omega_0) \geq p_0$, admits a rule-rational equilibrium in which players choose the coarsest partition and play a in all states.*

Proof. Let $\delta > 0$ be small enough such that: (1) the profiles (a, b_1, b_2) are δ-strict equilibria in game (A, u_0), (2) for each player i, $u_i(a) - u_i(b_i) > \delta$. Let $\epsilon > 0$ be large enough such that all games are ϵ-close to g^*: $\forall \omega \in \Omega$, $a^1 \in A^1$, $a^2 \in A^2$ $i \in \{1, 2\}$ $\left| u^i(a^1, a^2) - u_0^i(a^1, a^2) \right| \leq \epsilon$. Let $p_0 = \frac{\epsilon - \delta}{\epsilon}$. Then proposition 5 implies that there is a rule-rational equilibrium in which the players choose the coarsest partition, and play a in all states. $\qquad\square$

6 Stackelberg Stability and Act Rationality

This section shows that the existence of act-rational equilibria is closely related to the question of whether the Nash equilibria of the games in the environment are stable against the possibility that one of the players becomes a "Stackelberg" leader and commits in advance to a specific strategy. At the end of Section 2 we discussed the implications of the results of this section for the study of the value of information.

Given mixed action $x^i \in \Delta(A^i)$ of player i in a normal-form game $G = (A, u)$, let $BR(x^i) \subseteq A^j$ be the set of (pure) best replies of player j to x_i. Let $v^i(x^i) = \min_{x^j \in BR(x^i)} u_i(x^i, x^j)$ be the payoff that player i can guarantee himself when playing x^i before his opponent plays (under the assumption that his opponent best-replies to the player's commitment). Let $\bar{v}^i(G) = \max_{x^i \in \Delta(A^i)} v^i(x^i)$ be the payoff that player i can guarantee when playing as a Stackelberg leader (called a mixed-action Stackelberg payoff in Mailath & Samuelson (2006)). Let $NE(G)$ be the set of Nash equilibrium payoff vectors in G.

We say that a Nash equilibrium is Stackelberg-stable if no player can guarantee a higher payoff by being a Stackelberg leader. Formally:

Definition 9. A Nash equilibrium payoff vector $u \in NE(G)$ is *Stackelberg-stable* if for each player i, $\bar{v}_i(G) \leq u_i$.

A game is *Stackelberg-stable* if it admits at least one Stackelberg-stable equilibrium, and it is *Stackelberg-unstable* otherwise. Note that due to simple compactness arguments, each Stackelberg-unstable game G admits $\epsilon > 0$ such that $u_i < \bar{v}_i(G) - \epsilon$ for any Nash equilibrium payoff vector $u \in NE(G)$. The following fact shows a few simple families of Stackelberg-stable games. In particular, both constant-sum games and games with common interests have this property.

Fact 1. *Each of the following games are Stackelberg-stable.*

1. *Constant-sum games (all equilibria are stable).*

2. *Games with common interests (the Pareto-dominant equilibrium is stable).*

3. *Games in which both players have a strictly dominant action, e.g., Prisoner's Dilemma (the unique equilibrium is stable).*

The following propositions show the close relation between act-rational equilibria and Stackelberg-stable games. The first proposition shows that if all games are Stackelberg-stable, then an act-rational equilibrium exists.

Proposition 6. *If every game in the environment is Stackelberg-stable, then the environment admits an act-rational equilibrium.*

Proof. Consider the following strategy profile. Each player chooses the finest partition, and the players play on the equilibrium path a Stackelberg-stable equilibrium in each game. If any player deviates to a

different partition, then the players play a Nash equilibrium of the induced subgame, such that the deviating player earns weakly less than its payoff on the equilibrium path. The induced subgame admits such a Nash equilibrium because otherwise the deviating player would have a strictly profitable Stackelberg commitment in at least one game in the environment, and this would contradict our assumption that players play Stackelberg-stable equilibria in all states. The above argument implies that this strategy profile is an act-rational equilibrium.   □

The converse of Proposition 6 is false: environments with Stackelberg-unstable games may admit act-rational equilibria. Nevertheless, there is a general relationship between environments with act-rational equilibria and Stackelberg-stable games, which is studied in the following two propositions. The first proposition shows that a single Stackelberg-unstable game is enough to prevent the existence of an act-rational equilibrium in the environment.

The intuition of the result is that in a Stackelberg-unstable game there exists a player who would benefit if he had the ability to be a Stackelberg leader and commit an action before his opponent replies. Such a commitment is possible if the environment assigns sufficient probability to a zero-sum game in which this action is the unique minimax action. If this is the case, then the environment does not admit any act-rational equilibrium because one of the players would gain by coarsening his information, and thereby committing to play as if he were a Stackelberg leader. Formally:

Proposition 7. *For any Stackelberg-unstable game $G = (A, u_0)$, there exists an environment such that: (1) it includes only (A, u_0) and Stackelberg-stable games, and (2) it does not admit any act-rational equilibrium.*

Proof. For each player i, let $\bar{x}_i \in \Delta(A_i)$ be a (possibly mixed) action that guarantees a payoff of $\bar{v}^i(G)$ (by committing as a Stackelberg leader to play $s_i(v)$), and let (A, u_i) be a zero-sum game in which strategy $s_i(v)$ is the unique (possibly mixed) action that guarantees player i the value of the game (A, u_i). Let $\Omega = \{\omega_0, \omega_1, \omega_2\}$, and let $u(\omega_0) = u_0$, $u(\omega_1) = u_1$, and $u(\omega_2) = u_2$. We now show that if $p(\omega_0)$ is sufficiently small, then the environment (Ω, p, A, u) does not admit an act-rational equilibrium. Assume to the contrary that an act-rational equilibrium, τ, exists. The fact that game G is Stackelberg-unstable implies that for one of the players (say, player i): $u_i(x_\tau(\omega_0)) < \bar{v}_i(G) - \epsilon$. If $p(\omega_0)$ is sufficiently small then player i would strictly gain from deviating to the partition $\tilde{\pi}_i = \{\{\omega_0, \omega_i\}, \omega_j\}$. The argument for this is as follows. If $p(\omega_1)$ is sufficiently large with respect to $p(\omega_0)$ then it must be that in any equilibrium of the subgame after observing $(\tilde{\pi}_i, \bar{\pi}_j)$, player i's strategy is very close to $s_i(v)$ (because $s_i(v)$ is the unique minimax action in the zero-sum game (A, u_i)), and thus player i earns a payoff very close to $\bar{v}^i(G)$ in ω_0, which, for sufficiently small $p(\omega_0)$, is strictly larger than $u_i(x_\tau(\omega_0))$.   □

The second proposition shows that any act-rational equilibrium of a generic environment is preserved when a Stackelberg-stable game is added to the environment with small enough probability. Formally, a normal-form game is generic if all action profiles yield different payoffs, and an environment is generic if all its games are generic.

Proposition 8. *Let $((\Omega, p), A, u)$ be a generic environment that admits an act-rational equilibrium $\tau = ((\bar{\pi}, \sigma_1^*), (\bar{\pi}, \sigma_2^*))$. Extend the environment by adding an additional state $\tilde{\omega}$ to Ω with probability q in which a Stackelberg-stable game $(A, u(\tilde{\omega}))$ is played, and multiply the probabilities of any state $\omega \in \Omega$ by $1 - q$. Then if q is small enough, the new environment $((\Omega \cup \{\tilde{\omega}\}, \tilde{p}), A, u)$ admits an act-rational equilibrium that induces the same play in Ω.*

Proof. Consider the following strategy. Each player i chooses the finest partition in $\Omega \cup \{\tilde{\omega}\}$. The players play a Stackelberg-stable equilibrium in $\tilde{\omega}$ with payoff v, and follow (σ_1^*, σ_2^*) in Ω both on the equilibrium path and after any player has deviated in the first stage. If player i has deviated by choosing a partition in which $\tilde{\omega}$ is part of a non-trivial partition element, then the opponent plays in $\tilde{\omega}$ a best-reply to the deviator's mixed action (which is determined by his play in the remaining states in the partition element), which yields the deviator a payoff of at most v_i (such a best reply exists due to the Stackelberg stability of $\tilde{\omega}$). The fact that τ is act-rational and the environment is generic implies that quasi-strict equilibria are played in all states. This implies that for sufficiently small q, the above strategy profile is an act-rational equilibrium. $\qquad\square$

7   Feasible Partitions and Applications

In the basic model there were no restrictions on the set of feasible partitions. In this section we extend our model in order to deal with such restrictions, and we apply this extension to study lemon markets and partnership games.

7.1   Modeling Feasible Partitions

An environment with feasible partitions is an environment that also consists of a subset of allowed partitions to each player. Formally:

Definition 10. An *environment with feasible partitions* $((\Omega, p), A, u, (\Pi_1, \Pi_2))$ is a tuple where:

- $((\Omega, p), A, u)$ is an environment as defined in Section 3.

- For each player $i \in \{1, 2\}$, the set $\Pi_1 \subseteq \Pi_\Omega$ is a non-empty set of partitions, which is interpreted as the set of partitions that are feasible for player i.

The definition of the meta-game is adapted such that at stage 1, each player i chooses a partition from Π_i. Formally:

Definition 11. A *(behavior) strategy* for player i in the meta-game is a pair $\tau_i = (\mu_i, \sigma_i)$, where:

1. $\mu_i \in \Delta(\Pi_i)$ is a distribution over the set of feasible partitions.

2. $\sigma_i : \Pi_i \times \Pi_j \times \Omega \to \Delta(A^i)$ is a *playing rule,* i.e., a mapping that is measurable with respect to the partition of player i: $\sigma_i(\pi_i, \pi_j, \omega) = \sigma_i(\pi_i, \pi_j, \omega')$ whenever $\pi_i(\omega) = \pi_i(\omega')$.

As in the basic model, a rule-rational equilibrium is a subgame-perfect equilibrium of the meta-game.

We say that a profile of feasible partitions (Π_1, Π_2) is *regular* if each set Π_i admits a finest partition (denoted by $\bar{\pi}_i$); i.e., $\bar{\pi}_i$ is a refinement of all partitions in Π_i (note that $\bar{\pi}_i$ may differ from $\bar{\pi}$). Given an environment with regular feasible partitions, we say that a rule-rational equilibrium is *act-rational* if both players use only the finest partition.

We interpret the restriction to feasible partitions as a combination of three aspects of the strategic environment:

1. (Exogenous) incomplete information: It might be that some information about the opponent or about nature cannot be observed by the player. For example, in the lemon market application below, the seller cannot observe the private value of the seller's good.

2. Salient information: Some differences between states are so salient that the player always pays attention to them, and he cannot bundle together states that are so inherently different. For example, in the lemon market application below, the seller may bundle together two values of the good with a small difference between them (i.e., labeling both as having a "quite low" value), but he cannot bundle together a very low value and a very high value of his good.

3. Cultural factors: Cultural factors, such as the emotions, moral standards, and social norms discussed in the examples in the Introduction, allow one to bundle some groups of states together, but do not allow the bundling of other groups.

7.2   Application to Lemon Markets

Consider a seller who owns a car with a private value v that is uniformly distributed in the interval $\Omega = [0,1]$.[18] The buyer cannot observe this value, and thus has a single feasible partition (the coarsest one). The seller's feasible partitions bundle together "similar" values, under the interpretation that the seller only categorizes the good's value to several possible values, without paying attention to its exact value within the category. Formally, we say that that a seller's partition is feasible if each element in the partition is either a point or an interval with a length of at most $0 \leq d \leq 1$.

At the first stage of the meta-game the seller chooses his partition (interpreted as his rule to bundle together similar values). At the second stage the seller makes a take-it-or-leave-it offer to a buyer to purchase the car at price $0 \leq p$. The buyer either accepts or rejects the offer (while observing the seller's partition, but not the exact state or the specific partition element). The payoff (profit) of the seller is $p - v$ if there is a trade and 0 otherwise. The payoff to the buyer is $\alpha \cdot v - p$ if there is a trade and 0 otherwise, where $1 < \alpha < 2$ denotes the *surplus coefficient*: the good is worth α times more to the buyer than to the seller.

The case $d = 0$ is a classical lemon market (Akerlof, 1970; Samuelson & Bazerman, 1985), and it is well known that there is no trade in any equilibrium. The basic intuition is that the buyer infers from any price p that the value of the car for the seller is between zero and p (on average $\frac{p}{2}$), and thus its value for the buyer is on average $\alpha \cdot \frac{p}{2} < p$, and so the buyer never buys the good. Similarly, there cannot be trade in any act-rational equilibrium.

The following claim gives a minimal bound to the seller's payoff in any rule-rational equilibrium.

Claim 1. If $d > 0$ then the expected profit to the seller in any rule-rational equilibrium is at least $\frac{\alpha-1}{2} \cdot d^2$.

Proof. Assume to the contrary that there exists a rule-rational equilibrium in which the expected profit of the seller is less than $\frac{\alpha-1}{2} \cdot d^2$. Consider the following seller's deviation: choosing a partition that includes (as one of its elements) the interval $[0,d]$ and offering a price of $\alpha \cdot \frac{d}{2} - \epsilon$ (for any $\epsilon > 0$) when the value is in this interval. Observe that the buyer earns a positive payoff by accepting the offer, and thus he will accept it with probability one in any rule-rational equilibrium. Such a deviation yields the seller a payoff of at least $d \cdot \left(\alpha \cdot \frac{d}{2} - \epsilon - \frac{d}{2} \right) = \frac{\alpha-1}{2} \cdot d^2 - \epsilon \cdot d$ (because with probability d the seller sells a car with an expected value of $\frac{d}{2}$ at price $\alpha \cdot \frac{d}{2} - \epsilon$), and we get a contradiction (as the deviation is profitable for a sufficiently small $\epsilon > 0$ if the payoff in the original rule-rational equilibrium is less than $\frac{\alpha-1}{2} \cdot d^2$). $\square$

The above claim implies that for any $d > 0$, the environment does not admit any act-rational equilibrium, and that in any rule-rational equilibrium there is a positive level of trade.

---

[18] We slightly extend the formal model by allowing a continuum of states in this application and in the following one.

The following rule-rational equilibrium yields the seller a larger profit of $\frac{\alpha-1}{2\cdot(2-\alpha)^2}\cdot d^2$, as long as $d < 2 - \alpha$. The seller chooses a partition that includes (as one of its elements) the interval $\left[\frac{\alpha-1}{2-\alpha}\cdot d, \frac{1}{2-\alpha}\cdot d\right]$ and offers a price of $p = \frac{\alpha}{2\cdot(2-\alpha)}\cdot d$ if the value of the car is in this interval or lower (and he offers a higher price, which is always rejected, if the value of the car is higher than this interval). The buyer accepts this price (obtaining an expected payoff of zero because $p = \mathbf{E}\left(\alpha\cdot v|v \leq \frac{1}{2-\alpha}\cdot d\right)$) and rejects any higher price. The seller is indifferent between selling and not selling when the value of the good is inside the interval (because $p = \mathbf{E}\left(v|\frac{\alpha-1}{2-\alpha}\cdot d \leq v \leq \frac{1}{2-\alpha}\cdot d\right)$), and he earns $\frac{1}{(2-\alpha)}\cdot d$ when v is lower than this interval (because $\mathbf{E}\left(p - v|v \leq \frac{\alpha-1}{2-\alpha}\cdot d\right) = \frac{1}{2\cdot(2-\alpha)}\cdot d$). The probability of the latter event is $\frac{\alpha-1}{2-\alpha}\cdot d$, which implies that the expected profit is $\frac{\alpha-1}{2\cdot(2-\alpha)^2}\cdot d^2$.

It is relatively simple to see that this equilibrium induces the largest probability of trade and the largest profit to the seller among all the pure rule-rational equilibria (in which the buyer's choice of whether to accept trade is a deterministic function of the price), and we conjecture that it also holds for mixed rule-rational equilibria.

When $d \geq 2 - \alpha$, there is a rule-rational equilibrium in which the seller sells the car with probability one and obtains the total welfare gain from the efficient trade. The characterization of the rule-rational equilibrium is as follows. The seller chooses a partition that includes the interval $[\alpha - 1, 1]$, and offers a price of $p = \frac{\alpha}{2}$, regardless of the car's value. The buyer always accepts trade (and makes zero profit). The seller is indifferent between selling and not selling conditional on the value being in the interval (as its conditional expected value is $\frac{\alpha}{2}$), and earns all the positive gain from trade conditional on the value being lower than $\alpha - 1$.

7.3 Application to Team-Effort Games with Private Costs

Consider two partners that have simultaneously to choose whether to exert effort in the project (e) or not (n). The project succeeds if and only if both agents exert effort, in which case it yields each a payoff of one. Each player i may privately observe his own cost to exert effort, c_i, which is uniformly distributed in $[\delta, 1 + \delta]$ for some $\delta > 0$ (and we pay special interest in the case of a small δ). The total payoff for each player i is zero if he does not exert an effort, $1 - c_i$ if both players exert effort, and c_i if he is the sole player who exerts effort.

Observe that in a standard setup, the unique Nash equilibrium is the inefficient outcome in which no player ever exerts any effort. The sketch of the argument is as follows. The monotonicity of the payoffs with respect to the cost implies that both players use a threshold strategy: exert effort if and only if the private cost is below some threshold. Let x_1 (x_2) denote the threshold of player 1 (2). Assume without loss of generality that $x_j \leq x_i$ (i.e., player i uses a weakly larger threshold). Assume to the contrary that $x_i > 0$. When player i observes a private cost equal to the threshold ($c_i = x_i$) he is indifferent between exerting and not exerting effort:

$$0 = Pr\left(c_j \leq x_j\right) - x_i = x_j - \delta - x_i \leq -\delta < 0,$$

which yields a contradiction. Thus both players never exert effort.

The situation can be modeled as an environment with feasible partitions as follows. Let $\Omega = \Omega_1 \times \Omega_2 = [\delta, 1 + \delta] \times [\delta, 1 + \delta]$. Let the set of feasible partitions Π_i of each player i be the set of partitions in which all of their elements are Cartesian sets of the form $I_i \times \Omega_j$, where I_i is either a point or an interval of length at most d. That is, each player cannot observe his partner's private cost, and is able either to fully observe his own cost or to categorize it by bundling together similar costs with a difference of up

to d. Each player observes his own partition element, and the partner's partition (but not the partner's partition element or the exact state), and then chooses whether or not to exert effort.

The inefficient outcome in which both players never exert effort (regardless of their partition) is an act-rational equilibrium for any level of d. However, the following claim shows that a modest level of categorization, $d \geq 2 \cdot \delta$, is enough to sustain the efficient outcome in which both players always invest effort as a rule-rational equilibrium.

Claim 2. Let $d \geq 2 \cdot \delta$. Then there exists a rule-rational equilibrium in which both players exert effort regardless of their private costs.

Proof. Let the partition of each player i include the element $X_i = [1 - \delta, 1 + \delta]_i \times \Omega_j)$, and let each player i exert effort for all partition elements if the partner's partition includes the element X_j, and not exert effort otherwise. Observe that each player i is indifferent between exerting and not exerting effort conditional on observing the partition element X_i, and he strictly prefers to exert effort given any other signal. We interpret the partition element X_i as the player categorizing all high costs together, without paying attention to exactly how high the cost is. $\qquad \square$

8   Discussion

8.1   Empirical Implications

We believe our model and results make an important point concerning the interpretation of experimental results. Such results cannot be interpreted in isolation from subjects' experiences in strategic interactions outside the lab. When judged in isolation a certain behavior in the lab may seem grossly irrational, but when we embed the laboratory environment in a larger environment that includes also interactions outside the lab, such behavior may turn into a perfectly rational rule. Though this by no mean diminishes the importance of laboratory experiments in the social sciences, it does suggest that the discussion about whether subjects behave rationally or not in various games played in the lab has to be reconsidered.

Our model implies that the predictive power of a Nash equilibrium is expected to be stronger in Stackelberg-stable games, such as zero-sum games and coordination games, relative to Stackelberg-unstable games, in which the ability to commit and to play first is advantageous. This implications fits the stylized experimental facts (see Camerer, 2003, and the references there): (1) the aggregate play in zero-sum games is close to the Nash equilibrium (see also the empirical support for this from the behavior of professional tennis players in Walker & Wooders, 2001) ; (2) players are very good at coordinating on one of the pure equilibria of a coordination game; and (3) the predictive power of Nash equilibria in most other games is more limited.

Another insight from our the model is that the way in which people play a certain game depends not only on its payoff, but also on the propensity of playing the game within the relevant environment. Moreover, our model implies that games that are played proportionally more often should yield behavior that is closer to the predictions of Nash equilibrium.

8.2   Partial Observability

In this section, we sketch how the model and the results can be extended to a setup in which players only sometimes observe the partition of their opponent. Specifically, we change the definition of the two-stage meta-game that is induced by the environment by assuming that each player privately and independently

observes his opponent's partition with probability $0 \leq q \leq 1$ (à la Dekel et al. (2007)).[19] Accordingly, we extend the definition of a strategy to describe the behavior also after obtaining a non-informative signal about the opponent's partition. Finally, we define a rule-rational equilibrium as a sequential equilibrium (Kreps & Wilson, 1982) of the meta-game.

Observe that the case of $q = 0$ (i.e., non-observability of the opponent's partition) can only induce act-rational equilibria, and the case of $q = 1$ is equivalent to the basic model of Section 3. Most of our examples and results can be extended to this setup, as we sketch below.

Motivating Examples:

1. Example 1 (Prisoner's Dilemma and Stag Hunt): The efficient profile of choosing the coarse partition and playing (C, C) is a non-act-rational equilibrium for every $q \geq \frac{1}{35}$. This is so because, playing D is the unique best reply if one observes one's opponent deviating to the fine partition (note that the opponent is going to play D in the Prisoner's Dilemma). Thus a deviation to the fine partition yields a gain of $0.25 = 0.5 \cdot 0.5$ if undetected (due to playing D in the Prisoner's Dilemma), and a loss of $8.5 = 0.5 \cdot 8 + 0.5 \cdot 9$ if detected (due to the opponent playing D in both games). This implies that the deviation to the fine partition is unprofitable as long as $0.25 \cdot (1 - q) \leq 8.5 \cdot q \Leftrightarrow q \geq \frac{1}{35}$.

2. Example 2 (trade game): The efficient profile in which each player chooses the partition $\{\{\omega_1, \omega_2\}, \omega_3\}$, and plays T in $\{\omega_1, \omega_2\}$ (unless observing the opponent choosing the finest partition) is a non-act-rational equilibrium as long as $q \geq \frac{1}{4}$ because an undetected (detected) deviation to the finest partition yields a profit (loss) of $\frac{1}{3}$ ($\frac{4}{3}$); thus, the expected gain is negative as long as $q \geq \frac{1}{4}$.

3. Example 3 (with $p(\omega_1) = 0.5$, for concreteness) does not admit any act-rational equilibrium as long as $q \geq \frac{1}{8}$. If both players choose the fine partition then player 1 gains a payoff of one. If player 1 deviates to the coarse partition (and plays the dominant action C) he loses one utility point if the deviation is undetected, while he gains 8 utility points if the deviation is detected (and player 2 best replies to C by playing L). Similarly, the efficient profile in which $\pi_1^* = \pi_2^* = \underline{\pi}$, and the players play (C, L), is a non-act-rational equilibrium as long as $q \geq \frac{1}{8}$: a deviation of player 1 to the fine partition yields him a gain (loss) of one (eight) utility points if undetected (detected).

4. Example 4 (team-effort game): There exists a threshold $\bar{q} < 1$ such that the efficient rule-rational equilibrium remains an equilibrium for any $\bar{q} \leq q \leq 1$. This threshold depends on the distribution of αs, and it decreases as the probabilities assigned to lower values of α-s increase.

5. Example 5 (the ultimatum game): The non-act-rational equilibrium is no longer a rule-rational equilibrium when $q < 1$ (because deviating to the fine partition is a profitable deviation for player 2 for any $q < 1$). The adaptation of the example and its non-act-rational equilibrium to a setup with partial observability is left for future research.

6. Example 6 (chain-store game): The non-act-rational equilibrium in which the monopolist chooses the coarsest partition and fights all entrants is a rule-rational equilibrium for any q (because deviating to a finer partition does not yield any gain).

---

[19]The results can be extended also to a setup in which the signal about the opponent's partition is publicly observed (or has a positive probability of being observed by the opponent).

Environments with a uniform best-response correspondence.   Proposition 1 holds for any q (the proof essentially remains the same in the extended setup). We say that a correlated equilibrium is strictly undominated if for each player there exists a strict Nash equilibrium that yields the player a strictly worse payoff. One can adapt Proposition 2 so as to show that any strictly undominated correlated equilibrium can be implemented as a rule-rational equilibrium of a redundant equilibrium if $q < 1$ is sufficiently high.

Constant-sum environments and environments with common interests.   Propositions 3 and 4 hold for any q (the proofs essentially remain the same in the extended setup).

Environments with global rules.   One can slightly adapt the argument of Proposition 5 and show that it can also hold for a sufficiently high $q < 1$.

Stackelberg stability and act rationality.   Propositions 6 and 8 hold for any q (the arguments of the proof essentially remain the same in the extended setup). One can slightly adapt the argument of Proposition 7 and show that it can also hold for a sufficiently high $q < 1$.

It is also possible to combine the extension of partial observability with the extension of a restricted set of feasible partitions, and obtain similar results in the applications of Section 7 (lemon markets and partnership games) for a sufficiently high $q < 1$.

8.3   Value of Information

One can reinterpret an environment as an incomplete-information game in which each player has a minimal bound (the coarsest partition) and a maximal bound (the finest partition) on the amount of information he may acquire about the state of nature. Each player first chooses how much information to obtain within these bounds, and this choice is observable by the opponent, and then plays the incomplete information game induced by his choice. The partition profile of a rule-rational equilibrium thus describes an optimal information allocation (given his opponent's behavior): each player acquires an optimal amount of information, so that either obtaining more information or giving up information (weakly) reduces the player's payoff.

The results of Section 6 can be adapted to this interpretation as follows. Proposition 6 shows that if all games in the environment are Stackelberg-stable, then information has a positive value, and there exist an equilibrium in which both players choose to obtain the maximal amount of information. This extends related existing results for constant-sum games and games with common interests (both families of games are subsets of the set of Stackelberg-stable games). Proposition 7 shows that for any Stackelberg-unstable game, there is an environment in which it is the only game with this property, and information has a negative value: at least one of the players chooses to give up some information, and he will strictly lose by obtaining that additional information. Finally, Proposition 8 shows that if the value of information is positive at a specific equilibrium, then perturbing the environment by adding a Stackelberg-stable game with small probability maintains the positive value of information.

8.4   Future Research

We consider the current paper as a first step in a research agenda that studies the relations between rule rationality and environment-related commitments that operate on emotional levels (See also Winter et al., 2010). In what follows we briefly sketch two interesting directions to study in future research.

First, it would be interesting to extend the model by adding small costs for using more complex playing rules (i.e., finer partitions). Second, it would be interesting to study environments that include games in extensive form (rather than normal form), and adapt the notion of rule rationality also to actions at decision nodes, and not only to strategies as a whole (e.g., a rule by which your first move is always towards the terminal node with the highest payoff).

A   Correlated Equilibria and "Atomless" Environments

In this Appendix we show that the implementation of correlated equilibria by environments that only include copies of the same game does not depend on using "tailored" probability distributions. Specifically, we show that any undominated correlated equilibrium can be approximately induced by any large enough "atomless" environment (an environment in which each state has a small probability). Formally:

Definition 12. Given $\epsilon \geq 0$, a strategy profile $\tau^* = ((\mu_1^*, \sigma_1^*), (\mu_2^*, \sigma_2^*))$ is a *rule-rational ϵ-equilibrium* if for each player i:

1. Each playing-rule is an ϵ-best reply in each subgame: for each partition profile (π_1, π_2), player i, partition element $\pi_i(\omega)$, and action $a \in A^i$:

$$\sum_{\omega' \in \pi_i(\omega)} p(\omega') \cdot \left(u_i \left(\sigma_i^* (\pi_1, \pi_2 \omega'), \sigma_j^* (\pi_i, \pi_j, \omega') \right) - u_i \left(a, \sigma_j^* (\pi_i, \pi_j, \omega) \right) \right) \geq -\epsilon.$$

2. Each partition is an ϵ-best reply: for each chosen partition $\pi_i^* \in C(\mu_i^*)$ and for each partition π_i':

$$\sum_{\pi_j \in C(\mu_j^*)} \mu_j^* (\pi_j) \cdot \sum_{\omega \in \Omega} p(\omega) \cdot \left(u_i \left(\sigma_i^* (\pi_i^*, \pi_j, \omega), \sigma_j^* (\pi_i^*, \pi_j, \omega) \right) \right.$$
$$\left. - u_i \left(\sigma_i^* (\pi_i', \pi_j, \omega), \sigma_j^* (\pi_i', \pi_j, \omega) \right) \right) \quad \geq \quad -\epsilon.$$

Definition 13. Given a normal-form game (A, u) and $\epsilon \geq 0$, two correlated action profiles $x, y \in \Delta(A)$ are ϵ-close if $\max_{a \in A} (|x(a) - y(a)|) \leq \epsilon$.

Definition 14. Given a normal-form game (A, u_0), let $\bar{u}_0 = \max_{a, a' \in A, i \in \{1,2\}} |u_i(a) - u_i(a')|$ be the maximal difference between any two values in the payoff matrix.

Proposition 9. *Let (A, u_0) be a game, let ζ be an undominated correlated equilibrium in (A, u_0), and let $\epsilon > 0$. Let $((\Omega, p), A, u_0)$ be an environment of strategically equivalent situations that satisfies $p(\omega) < \epsilon/\bar{u}_0 \; \forall \omega \in \Omega$. Then, the environment $((\Omega, p), A, u_0)$ admits a rule-rational ϵ-equilibrium that induces an ex-ante outcome, which is ϵ-close to ζ.*

Proof. Let $x_\zeta \in \Delta(A)$ be the outcome of the correlated equilibrium ζ. It is immediate that there exists a function $g : \Omega \to A$ that identifies each set of states with an action profile in A, such that for each action profile $a \in A$, $|p(\{\omega | g(\omega) = a\}) - x_\zeta(a)| < \epsilon/\bar{u}_0$. Let $g_i(\omega) : \Omega \to A_i$ be the function that identifies each state with the induced action of player i (i.e, $g(\omega) = (g_1(\omega), g_2(\omega))$). For each player i let partition π_i be defined as follows: $\pi_i(\omega_0) = \{\omega | g_i(\omega) = g_i(\omega_0)\}$. Consider the following strategy profile $\tau_{\zeta, \epsilon} = ((\pi_1, \sigma_1), (\pi_2, \sigma_2))$: on the equilibrium path $\sigma_i(\pi_1, \pi_2, \omega) = g_i(\omega)$, and if player i deviates and chooses a different partition, then both agents play in all states the Nash equilibrium that yields player i a worse payoff. It is immediate to see that $\tau_{\zeta, \epsilon}$ is a rule-rational ϵ-equilibrium that yields an outcome that is $\epsilon/\bar{u}_0$ close to ζ. $\square$

References

Akerlof, G. A. (1970). The market for "lemons": Quality uncertainty and the market mechanism. *The Quarterly Journal of Economics*, 84(3), 488–500.

Aumann, R. J. (1974). Subjectivity and correlation in randomized strategies. *Journal of Mathematical Economics*, 1(1), 67–96.

Aumann, R. J. (1987). Correlated equilibrium as an expression of Bayesian rationality. *Econometrica*, 55(1), 1–18.

Aumann, R. J. (2008). Rule rationality vs. act rationality. *The Hebrew University, Center for the Study of Rationality*, DP-497.

Aumann, R. J. & Sorin, S. (1989). Cooperation and bounded recall. *Games and Economic Behavior*, 1(1), 5–39.

Azrieli, Y. (2009). Categorizing others in a large game. *Games and Economic Behavior*, 67(2), 351–362.

Bade, S., Haeringer, G., & Renou, L. (2009). Bilateral commitment. *Journal of Economic Theory*, 144(4), 1817–1831.

Bassan, B., Gossner, O., Scarsini, M., & Zamir, S. (2003). Positive value of information in games. *International Journal of Game Theory*, 32(1), 17–31.

Baumol, W. J. & Quandt, R. E. (1964). Rules of thumb and optimally imperfect decisions. *The American Economic Review*, 54(2), 23–46.

Bednar, J., Chen, Y., Liu, T. X., & Page, S. (2012). Behavioral spillovers and cognitive load in multiple games: An experimental study. *Games and Economic Behavior*, 74(1), 12–31.

Ben-Porath, E. (2003). Cheap talk in games with incomplete information. *Journal of Economic Theory*, 108(1), 45–71.

Blackwell, D. (1953). Equivalent comparisons of experiments. *The Annals of Mathematical Statistics*, 24(2), 265–272.

Bomze, I. M. (1986). Non-cooperative two-person games in biology: A classification. *International Journal of Game Theory*, 15(1), 31–57.

Camerer, C. (2003). *Behavioral Game Theory: Experiments in Strategic Interaction*. Princeton University Press: Princeton, NJ.

Dekel, E., Ely, J. C., & Yilankaya, O. (2007). Evolution of preferences. *Review of Economic Studies*, 74(3), 685–704.

Ellison, G. & Fudenberg, D. (1993). Rules of thumb for social learning. *Journal of Political Economy*, 101, 612–643.

Ellison, G. & Holden, R. (2013). A theory of rule development. *Journal of Law, Economics and Organization*, forthcoming.

Evangelista, F. S. & Raghavan, T. (1996). A note on correlated equilibrium. *International Journal of Game Theory*, 25(1), 35–41.

Fershtman, C., Judd, K. L., & Kalai, E. (1991). Observable contracts: Strategic delegation and cooperation. *International Economic Review*, 32(3), 551–559.

Frenkel, S., Heller, Y., & Teper, R. (2014). The endowment effect as a blessing. mimeo. https://sites.google.com/site/yuval26/endowment-blessing.pdf.

Gale, J., Binmore, K. G., & Samuelson, L. (1995). Learning to be imperfect: The ultimatum game. *Games and Economic Behavior*, 8(1), 56–90.

Gossner, O. (2010). Ability and knowledge. *Games and Economic Behavior*, 69(1), 95–106.

Govindan, S. & Wilson, R. (2008). On forward induction. *Econometrica*, 77(1), 1–28.

Grimm, V. & Mengel, F. (2012). An experiment on learning in a multiple games environment. *Journal of Economic Theory*, 147(6), 2220–2259.

Güth, W. & Kocher, M. G. (2014). More than thirty years of ultimatum bargaining experiments: Motives, variations, and a survey of the recent literature. *Journal of Economic Behavior & Organization*, 108, 396–409.

Guth, W. & Yaari, M. (1992). Explaining reciprocal behavior in simple strategic games: An evolutionary approach. In U. Witt (Ed.), *Explaining Process and Change: Approaches to Evolutionary Economics*. University of Michigan Press, Ann Arbor.

Harsanyi, J. C. (1977). Rule utilitarianism and decision theory. *Erkenntnis*, 11(1), 25–53.

Hart, S. (2002). Evolutionary dynamics and backward induction. *Games and Economic Behavior*, 41(2), 227–264.

Hart, S. & Mas-Colell, A. (2000). A simple adaptive procedure leading to correlated equilibrium. *Econometrica*, 68(5), 1127–1150.

Heifetz, A., Shannon, C., & Spiegel, Y. (2007). What to maximize if you must. *Journal of Economic Theory*, 133(1), 31–57.

Heller, Y. (2014). Stability and trembles in extensive-form games. *Games and Economic Behavior*, 84, 132–136.

Heller, Y. (2015). Three steps ahead. *Theoretical Economics*, 10, 203–241.

Huck, S., Jehiel, P., & Rutter, T. (2011). Feedback spillover and analogy-based expectations: A multi-game experiment. *Games and Economic Behavior*, 71(2), 351–365.

Israel, S., Hart, E., & Winter, E. (2014). Oxytocin decreases accuracy in the perception of social deception. *Psychological Science*, 25(1), 293–295.

Jehiel, P. (2005). Analogy-based expectation equilibrium. *Journal of Economic Theory*, 123(2), 81–104.

Jehiel, P. & Koessler, F. (2008). Revisiting games of incomplete information with analogy-based expectations. *Games and Economic Behavior*, 62(2), 533–557.

Jung, Y. J., Kagel, J. H., & Levin, D. (1994). On the existence of predatory pricing: An experimental study of reputation and entry deterrence in the chain-store game. *RAND Journal of Economics*, 25(1), 72–93.

Kamien, M. I., Tauman, Y., & Zamir, S. (1990). On the value of information in a strategic conflict. *Games and Economic Behavior*, 2(2), 129–153.

Kreps, D. M. & Wilson, R. (1982). Sequential equilibria. *Econometrica*, 50(4), 863–894.

Lehrer, E., Rosenberg, D., & Shmaya, E. (2010). Signaling and mediation in games with common interests. *Games and Economic Behavior*, 68(2), 670–682.

Mailath, G. J. & Samuelson, L. (2006). *Repeated Games and Reputations.* Oxford university press, Oxford.

Maynard-Smith, J. & Price, G. (1973). The logic of animal conflict. *Nature*, 246, 15.

Mengel, F. (2012). Learning across games. *Games and Economic Behavior*, 74(2), 601–619.

Mohlin, E. (2011). Optimal categorization. *Stockholm School of Economics Working Paper Series in Economics and Finance.*

Morris, S. & Ui, T. (2004). Best response equivalence. *Games and Economic Behavior*, 49(2), 260–287.

Moulin, H. & Vial, J. (1978). Strategically zero-sum games: The class of games whose completely mixed equilibria cannot be improved upon. *International Journal of Game Theory*, 7(3), 201–221.

Nachbar, J. H. (1990). Evolutionary selection dynamics in games: Convergence and limit properties. *International Journal of Game Theory*, 19(1), 59–89.

Neral, J. & Ochs, J. (1992). The sequential equilibrium theory of reputation building: A further test. *Econometrica*, 60(5), 1151–1169.

Ochs, J. & Roth, A. E. (1989). An experimental study of sequential bargaining. *The American Economic Review*, 79(3), 355–384.

Osborne, M. J. & Rubinstein, A. (1994). *A Course in Game Theory.* MIT Press, Cambrdige, MA.

Renou, L. (2009). Commitment games. *Games and Economic Behavior*, 66(1), 488–505.

Roth, A. E., Prasnikar, V., Okuno-Fujiwara, M., & Zamir, S. (1991). Bargaining and market behavior in Jerusalem, Ljubljana, Pittsburgh, and Tokyo: An experimental study. *The American Economic Review*, 81(5), 1068–1095.

Samuelson, L. (2001). Analogies, adaptation, and anomalies. *Journal of Economic Theory*, 97(2), 320–366.

Samuelson, W. & Bazerman, M. H. (1985). The winner's curse in bilateral negotiations. In V. Smith (Ed.), *Research in Experimental Economics*, volume III. JAI Press: Greenwich, CT.

Schelling, T. C. (1960). *The Strategy of Conflict.* Harvard University Press, Cambridge, MA.

Selten, R. (1975). Reexamination of the perfectness concept for equilibrium points in extensive games. *International Journal of Game Theory*, 4(1), 25–55.

Selten, R. (1978). The chain store paradox. *Theory and Decision*, 9(2), 127–159.

Selten, R. (1983). Evolutionary stability in extensive two-person games. *Mathematical Social Sciences*, 5(3), 269–363.

Taylor, P. & Jonker, L. (1978). Evolutionary stable strategies and game dynamics. *Mathematical Biosciences*, 40(1), 145–156.

Walker, M. & Wooders, J. (2001). Minimax play at Wimbledon. *The American Economic Review*, 91(5), 1521–1538.

Weibull, J. (1997). *Evolutionary Game Theory*. MIT Press, Cambridge, MA.

Winter, E., Garcia-Jurado, I., & Mendez-Naya, L. (2010). Mental equilibrium and rational emotions. *The Hebrew University, Center for the Study of Rationality, DP-521.*